



# **ONTAP 9.14.1 EMS reference**

## **ONTAP EMS reference**

NetApp  
November 25, 2024

# Table of Contents

- ONTAP 9.14.1 EMS reference . . . . . 1
- accesscache events . . . . . 2
  - accesscache.cfg events . . . . . 2
  - accesscache.nearlimits events . . . . . 2
  - accesscache.reachedlimits events . . . . . 3
- acp events . . . . . 4
  - acp.auto events . . . . . 4
  - acp.bind events . . . . . 4
  - acp.commerr events . . . . . 4
  - acp.comherrpersist events . . . . . 5
  - acp.dev events . . . . . 5
  - acp.downrev events . . . . . 6
  - acp.exp events . . . . . 6
  - acp.fw events . . . . . 9
  - acp.init events . . . . . 10
  - acp.invalid events . . . . . 12
  - acp.ioxm events . . . . . 12
  - acp.locked events . . . . . 13
  - acp.need events . . . . . 14
  - acp.no events . . . . . 15
  - acp.post events . . . . . 15
  - acp.shelf events . . . . . 16
  - acp.sock events . . . . . 19
  - acp.update events . . . . . 19
  - acp.updateconfiguration events . . . . . 20
  - acp.upgrade events . . . . . 20
  - acp.voltage events . . . . . 21
- adminapi events . . . . . 22
  - adminapi.input events . . . . . 22
  - adminapi.xml events . . . . . 22
- adt events . . . . . 24
  - adt.aggrcb events . . . . . 24
  - adt.aucb events . . . . . 24
  - adt.cmdq events . . . . . 25
  - adt.cmdthrd events . . . . . 25
  - adt.dest events . . . . . 25
  - adt.max events . . . . . 27
  - adt.nfs events . . . . . 28
  - adt.service events . . . . . 28
  - adt.stgvol events . . . . . 30
  - adt.wrkrq events . . . . . 31
  - adt.wrkrthrd events . . . . . 32
- aggr events . . . . . 33

aggr.full events	33
aggr.grow events	33
aggr.inconsistent events	33
aggr.max events	34
aggr.provisioning events	34
aggr.relocation events	35
aggr.resize events	35
aggr.wafliron events	37
aggrcopy events	38
aggrcopy.dst events	38
aggrcopy.restart events	40
aggrcopy.src events	40
air events	41
air.certify events	41
air.cleared events	41
air.corrected events	42
air.disable events	42
air.enabled events	43
air.evicted events	43
air.fc events	44
air.qtm events	44
air.removed events	45
air.repair events	45
air.repaired events	46
air.vplus events	64
am events	65
am.dest events	65
api events	66
api.engine events	66
api.fileio events	66
api.input events	67
api.output events	67
app events	69
app.cg events	69
app.log events	69
app.temp events	71
application events	75
application.volume events	75
arl events	76
arl.abort events	76
arl.aggrdesttimeout events	76
arl.aggronlinefailed events	77
arl.aggronlinetimeout events	77
arl.aggrstart events	77
arl.connecttodestfailed events	78

arl.netra events	78
arl.opaborted events	81
arl.opfinished events	82
arl.opstart events	82
arl.postmignotok events	83
arl.precmt events	83
arl.subsystemabort events	86
arl.veto events	87
arw events	93
arw.analytics events	93
arw.auto events	95
arw.new events	95
arw.snapshot events	96
arw.volume events	96
arw.vserver events	97
asa events	98
asa.personality events	98
asup events	100
asup.aods events	100
asup.config events	100
asup.general events	101
asup.post events	103
asup.smtp events	103
asup.throttle events	104
asynclcli events	105
asynclcli.message events	105
asynclcli.permscheck events	105
asyncldelete events	107
asyncldelete.cancel events	107
asyncldelete.message events	107
ata events	109
ata.excessive events	109
ata.timeout events	109
audit events	111
audit.cap events	111
audit.log events	111
auditlog events	114
auditlog.change events	114
auto events	115
auto.cds events	115
auto.policy events	115
auto.update events	117
aws events	121
aws.dev events	121
aws.instance events	122

aws.nic events	122
backup events	124
backup.abort events	124
backup.fail events	124
backup.flip events	126
bcomd events	127
bcomd.fcvmelifadded events	127
bcomd.fcvmelifmovefail events	127
bcomd.fcvmelifremoved events	128
bcomd.fcplifadded events	128
bcomd.fcplifremoved events	128
bcomd.igroup events	129
bcomd.lun events	129
bcomd.vdisk events	130
bcomd.volume events	131
bdfu events	132
bdfu.selected events	132
block events	133
block.giveback events	133
boot events	134
boot.bootmenu events	134
boot.shutdown events	134
boot.varfs events	134
boot.verify events	137
bootfs events	138
bootfs.env events	138
bridge events	139
bridge.discovered events	139
bridge.duplicate events	139
bridge.removed events	140
bridge.sas events	140
bridge.timestamp events	141
bsdsocket events	142
bsdsocket.queuelimit events	142
bucket events	143
bucket.soft events	143
cache events	144
cache.ttl events	144
cachemirror events	145
cachemirror.offline events	145
callhome events	146
callhome.aggr events	146
callhome.andu events	148
callhome.appdm events	149
callhome.arw events	149

callhome.bad events	150
callhome.battery events	150
callhome.bcom events	153
callhome.bdfu events	153
callhome.bootfs events	154
callhome.bridge events	154
callhome.c events	155
callhome.carrier events	157
callhome.ch events	158
callhome.chassis events	160
callhome.checksum events	165
callhome.clam events	165
callhome.client events	166
callhome.clus events	168
callhome.config events	170
callhome.coredump events	171
callhome.corefile events	171
callhome.cpu events	172
callhome.css events	173
callhome.ctran events	173
callhome.dafs events	174
callhome.dbs events	175
callhome.device events	190
callhome.disk events	190
callhome.diskown events	192
callhome.disks events	192
callhome.dr events	193
callhome.dsk events	193
callhome.dumptriage events	202
callhome.encl events	203
callhome.ext events	203
callhome.fans events	204
callhome.fc events	204
callhome.fcp events	205
callhome.fdisk events	207
callhome.flash events	210
callhome.fs events	211
callhome.hainterconnect events	212
callhome.hba events	212
callhome.hm events	213
callhome.hostport events	214
callhome.hwaccel events	215
callhome.ibretimerprog events	216
callhome.invoke events	216
callhome.iscsi events	219

callhome.isp events	219
callhome.jpc events	220
callhome.kernel events	220
callhome.link events	222
callhome.management events	222
callhome.mcc events	223
callhome.mccip events	232
callhome.mdb events	233
callhome.mhost events	234
callhome.micro events	235
callhome.mlm events	236
callhome.multi events	237
callhome.ndo events	237
callhome.netif events	243
callhome.netinet events	244
callhome.nht events	245
callhome.no events	245
callhome.nse events	246
callhome.nvdimn events	246
callhome.nve events	247
callhome.nvme events	248
callhome.nvmem events	249
callhome.nvram events	250
callhome.okm events	251
callhome.ondemand events	252
callhome.partner events	254
callhome.performance events	254
callhome.pool events	256
callhome.raid events	257
callhome.raidtree events	258
callhome.reboot events	259
callhome.rnode events	266
callhome.root events	266
callhome.rtc events	267
callhome.sas events	267
callhome.sb events	268
callhome.sblade events	269
callhome.scsi events	273
callhome.sdisk events	273
callhome.ses events	276
callhome.sfo events	277
callhome.shlf events	282
callhome.shutdown events	285
callhome.sm events	285
callhome.snapmirror events	286

callhome.sp events	287
callhome.spares events	288
callhome.spm events	288
callhome.stackmon events	289
callhome.svm events	289
callhome.svo events	293
callhome.sync events	294
callhome.syncm events	295
callhome.syncsm events	297
callhome.sys events	297
callhome.tgt events	298
callhome.unsupported events	299
callhome.ups events	299
callhome.user events	302
callhome.vdom events	302
callhome.vfiler events	303
callhome.vol events	304
callhome.wafl events	306
callhome.wafliron events	308
callhome.weekly events	308
cam events	310
cam.dev events	310
cam.timeout events	310
cap events	312
cap.cache events	312
ccma events	314
ccma.disabled events	314
ccma.perfstatd events	314
ccma.quota events	315
cde events	317
cde.undo events	317
cdpd events	318
cdpd.recv events	318
cecc_log events	319
cecc_log.dropped events	319
cecc_log.entry events	319
cecc_log.entry_no_syslog events	319
cecc_log.masked events	320
cecc_log.summary events	320
certificate events	322
certificate.insecure events	322
certificate.invalid events	323
certificate.node events	323
cf events	324
cf.autogbcanceled events	324



cf.badshelfcountmsg events . . . . .	324
cf.badshelfcountverson events . . . . .	324
cf.boot events . . . . .	325
cf.cdo events . . . . .	325
cf.disk events . . . . .	326
cf.diskinventory events . . . . .	331
cf.fm events . . . . .	332
cf.fmmbx events . . . . .	360
cf.fmns events . . . . .	361
cf.fsm events . . . . .	361
cf.giveback events . . . . .	383
cf.headswap events . . . . .	385
cf.hwassist events . . . . .	387
cf.ic events . . . . .	397
cf.mccip events . . . . .	404
cf.misc events . . . . .	405
cf.mode events . . . . .	410
cf.multidisk events . . . . .	412
cf.nducomplete events . . . . .	413
cf.nduincomplete events . . . . .	413
cf.ndustart events . . . . .	414
cf.nm events . . . . .	414
cf.noautogb events . . . . .	415
cf.nopartner_takeover events . . . . .	415
cf.nullshelfcountmsg events . . . . .	416
cf.partner events . . . . .	416
cf.reservation events . . . . .	418
cf.rsrc events . . . . .	419
cf.rv events . . . . .	423
cf.shutdown events . . . . .	425
cf.sysid events . . . . .	426
cf.takeover events . . . . .	426
cf.transition events . . . . .	429
cifs events . . . . .	430
cifs.does events . . . . .	430
cifs.domainpwd events . . . . .	430
cifs.lug events . . . . .	431
cifs.replay events . . . . .	432
cifs.server events . . . . .	432
cifs.shadowcopy events . . . . .	433
cifs.static events . . . . .	434
clam events . . . . .	435
clam.disable events . . . . .	435
clam.enable events . . . . .	435
clam.failover events . . . . .	436

clam.heartbeat events	437
clam.invalid events	437
clam.mailbox events	438
clam.node events	438
clam.partner events	441
clam.peer events	441
clam.reenable events	442
clam.start events	442
clam.stop events	443
clam.takeover events	443
clam.update events	444
clam.valid events	445
clam.voting events	445
clone events	447
clone.split events	447
cloud events	448
cloud.aws events	448
cloud.azurecredsinvalid events	451
cloud.bin events	451
cloud.creds events	452
cloud.gpccredsinvalid events	453
clus events	454
clus.net events	454
clus.reboot events	455
cluster events	457
cluster.epsilon events	457
cluster.ha events	458
cluster.image events	458
cluster.network events	459
cluster.op events	459
cluster.port events	462
cluster.repl events	463
cluster.switch events	464
cluster.vol events	467
cmd events	470
cmd.result events	470
cmd.sasadmin events	470
cmd.shelfchk events	470
cmd.start events	471
cmds events	473
cmds.config events	473
cmds.software events	473
cmds.sysconf events	474
cmpr events	478
cmpr.algo events	478

- cna events . . . . . 479
  - cna.dump events . . . . . 479
  - cna.initfailed events . . . . . 479
  - cna.sfp events . . . . . 480
- command events . . . . . 481
  - command.complex events . . . . . 481
- config events . . . . . 482
  - config.atanotsupported events . . . . . 482
  - config.badlocalswitch events . . . . . 482
  - config.badpoolassign events . . . . . 482
  - config.badslot events . . . . . 483
  - config.badswitchlocal events . . . . . 483
  - config.bridge events . . . . . 484
  - config.cfobothabmodule events . . . . . 486
  - config.cfomixedabmodule events . . . . . 487
  - config.differentswitchport events . . . . . 487
  - config.diffpool events . . . . . 488
  - config.disknoswitch events . . . . . 488
  - config.exceededha events . . . . . 488
  - config.fcsameloop events . . . . . 489
  - config.habadbank1 events . . . . . 489
  - config.habadbank2 events . . . . . 490
  - config.habadremote events . . . . . 490
  - config.invalid events . . . . . 490
  - config.localswitch events . . . . . 491
  - config.mirroring events . . . . . 491
  - config.mirrornotmultipath events . . . . . 492
  - config.mirrorscsidisk events . . . . . 492
  - config.mixedabmodule events . . . . . 493
  - config.more2path events . . . . . 493
  - config.move events . . . . . 494
  - config.multipath events . . . . . 494
  - config.noadisk events . . . . . 494
  - config.nobconnect events . . . . . 495
  - config.nocheck events . . . . . 495
  - config.nomultipath events . . . . . 496
  - config.nopartnerdisks events . . . . . 496
  - config.nopartnerluns events . . . . . 497
  - config.noswitchcount events . . . . . 497
  - config.noswitchcountfordisk events . . . . . 497
  - config.notmultipath events . . . . . 498
  - config.oneswitch events . . . . . 498
  - config.pvlan events . . . . . 499
  - config.repl events . . . . . 499
  - config.sameha events . . . . . 500

config.sameslot events	501
config.sameswitch events	501
config.sameswitchwarn events	502
config.sas2domain events	502
config.scsidisk events	502
configbr events	504
configbr.backupcompleted events	504
configbr.backupfailed events	504
configbr.deletefailed events	505
configbr.distributefailed events	505
configbr.initfailed events	505
configbr.interror events	506
configbr.nodesturl events	506
configbr.snapshotdelfailed events	507
configbr.unpubfailed events	507
configbr.uploadfailed events	508
connectx events	509
connectx.cmdinterface events	509
connectx.fatalerror events	509
connectx.ibcabledetected events	509
connectx.iblinkretrainnotreqd events	510
connectx.iblinkretrainreqd events	510
connectx.ibqsfpdumpctrl events	511
connectx.ibqsfpdumpdata events	511
connectx.ibretimerfailcount events	512
connectx.ibretimerprogrmpass events	512
connectx.ibretimerskipprogrm events	513
connectx.ibunsupportcable events	513
connectx.portdisabled events	513
copy events	515
copy.dst events	515
coredump events	516
coredump.all events	516
coredump.compress events	516
coredump.compression events	516
coredump.coreid events	517
coredump.delete events	517
coredump.disk events	518
coredump.dr events	518
coredump.dump events	520
coredump.encryption events	527
coredump.findcore events	528
coredump.host events	530
coredump.job events	531
coredump.long events	532

coredump.mcc events	533
coredump.micro events	533
coredump.nosavecore events	534
coredump.nvlog events	534
coredump.save events	535
coredump.segment events	540
coredump.shutdown events	541
coredump.spare events	541
coredump.sparecore events	543
coredump.sparecoreinit events	544
cpeer events	545
cpeer.access events	545
cpeer.addr events	545
cpeer.authentication events	547
cpeer.create events	548
cpeer.dns events	549
cpeer.duplicate events	550
cpeer.ip events	550
cpeer.localid events	551
cpeer.long events	551
cpeer.modify events	552
cpeer.out events	552
cpeer.passphrase events	553
cpeer.ping events	553
cpeer.pm events	554
cpeer.pp events	554
cpeer.psk events	555
cpeer.receiver events	557
cpeer.recv events	557
cpeer.unavailable events	558
cpeer.xcm events	558
cr events	560
cr.data events	560
cr.del events	560
cr.fix events	562
cr.rdb events	565
cr.redir events	566
cr.snapshot events	566
cr.sync events	567
crypto events	568
crypto.debug events	568
crypto.export events	568
crypto.import events	568
crypto.key events	569
crypto.okmrecovery events	570

crypto.ssal events	570
cs events	572
cs.create events	572
cs.endpoint events	572
csm events	573
csm.badauthaccess events	573
csm.badconnection events	573
csm.badplainaccess events	574
csm.ctfallbackactiveopen events	574
csm.ctfallbackswitchless events	575
csm.homeportinfounexpected events	575
csm.mismatchremotedevice events	576
csm.nomemory events	576
csm.rdmarecoveractiveopen events	577
csm.rdmarecoverswitchless events	577
csm.unmarshallerror events	577
ctl events	579
ctlctlbe events	579
ctl.session events	580
ctran events	582
ctran.disabled events	582
ctran.enabled events	582
ctran.gbuf events	582
ctran.gc events	583
ctran.group events	584
ctran.jpc events	584
ctran.net events	586
ctrl events	588
ctrl.link events	588
ctrl.rdma events	589
daasync events	591
daasync.capacity events	591
daasync.inquiry events	591
daasync.periph events	591
daregister events	593
daregister.periph events	593
db events	594
db.api events	594
db.assignments events	594
db.block events	595
db.bulk events	596
db.capacity events	596
db.clone events	597
db.cluster events	597
db.db events	599

dbns.director events	600
dbns.discon events	601
dbns.disconn events	601
dbns.drive events	602
dbns.drives events	603
dbns.ekm events	604
dbns.encr events	605
dbns.ensemble events	606
dbns.exception events	607
dbns.fault events	607
dbns.file events	608
dbns.generic events	609
dbns.ie events	610
dbns.ikm events	611
dbns.install events	612
dbns.limit events	612
dbns.maintenance events	613
dbns.master events	614
dbns.mem events	614
dbns.network events	616
dbns.node events	617
dbns.platform_hardware events	619
dbns.prestartup events	620
dbns.provisioned events	620
dbns.raid events	621
dbns.remote events	622
dbns.remrep events	623
dbns.schedule events	625
dbns.secondary events	625
dbns.service events	626
dbns.slice events	627
dbns.snapshot events	628
dbns.snmp events	629
dbns.software events	630
dbns.ssl events	630
dbns.stat events	631
dbns.sw events	631
dbns.unexpected events	632
dbns.unresponsive events	633
dbns.upgrade events	634
dbns.valence events	634
dbns.volumes events	635
dcns events	637
dcns.framework events	637
dev events	638

dev.driver events	638
dev.path events	638
device events	639
device.install events	639
device.latency events	639
device.max events	640
dfu events	641
dfu.badbridgeshelfconfig events	641
dfu.clusterdisknotready events	641
dfu.clustericdown events	642
dfu.clustericerror events	642
dfu.clusterpartner events	643
dfu.clustertakeover events	643
dfu.disksizechanged events	644
dfu.firmwaredownloading events	644
dfu.firmwaredownrev events	645
dfu.firmwareuptodate events	645
dfu.fwdownloadfailedmultfiles events	646
dfu.fwdownloadfailedstatus events	646
dfu.fwdownloadfailedwbuf events	647
dfu.init events	647
dfu.mallocfailed events	648
dfu.nomemory events	648
dfu.readtestfailed events	649
dfu.readtestfailedstatus events	649
dfu.refreshinquiryfailed events	650
dfu.suspended_disasterrecovery events	650
diff events	652
diff.svcs events	652
dipmon events	653
dipmon.dataaip events	653
director events	654
director.aggr events	654
director.dataaip events	654
discover events	656
discover.6500 events	656
discoveryd events	657
discoveryd.search events	657
disk events	658
disk.adaptererror events	658
disk.bdfu events	658
disk.capacity events	659
disk.checksum events	659
disk.cksum events	660
disk.ddr events	660



disk.diagnostic events	664
disk.duplicate events	664
disk.dynamicqual events	665
disk.encrypt events	667
disk.encryptaccesserr events	668
disk.encryptaccessfwerr events	669
disk.encryptcmdfailed events	669
disk.encryptfipspassphnosync events	670
disk.encryptgetmsidfailed events	670
disk.encryptnosupport events	671
disk.encryptprotdiff events	671
disk.encryptsupportreqd events	672
disk.eos events	672
disk.exceed events	673
disk.fail events	673
disk.failmsg events	674
disk.fakereassignsuccess events	674
disk.fli events	675
disk.fw events	676
disk.healthtrigger events	677
disk.init events	677
disk.iomediumerror events	683
disk.ioereassignfailed events	683
disk.ioereassignsuccess events	684
disk.ioerecoverederror events	684
disk.lipstormdetect events	686
disk.lostwritedetected events	687
disk.lun events	687
disk.maint events	689
disk.max events	690
disk.min events	691
disk.outofservice events	691
disk.partdisknotsuppmcc events	691
disk.partition events	692
disk.partner events	692
disk.readreservationfailed events	693
disk.releasefailed events	694
disk.reservedelay events	695
disk.reserveerror events	695
disk.reservefailed events	696
disk.rewritedatafailed events	696
disk.sanit events	697
disk.selectiontimeout events	701
disk.senseerror events	701
disk.senseinfontotvalid events	702

disk.setarraylqd events	702
disk.sizegrew events	702
disk.sizeshrank events	703
disk.spindleerror events	703
disk.stalledtoolong events	704
disk.timeout events	705
disk.toomanynvmespindles events	706
disk.toomanypartitioneddskis events	706
disk.toomanyreassignments events	706
disk.toomanyspindles events	707
disk.uid events	707
disk.unabletorewritedata events	708
disk.vendorspecificcode events	708
disk.writeverifyerror events	709
diskown events	710
diskown.aashelf events	710
diskown.assign events	712
diskown.assignexceeded events	713
diskown.autoassign events	713
diskown.autoassignproblem events	714
diskown.bayctrlr events	715
diskown.changingowner events	716
diskown.cleaningownership events	716
diskown.completion events	717
diskown.delayed events	717
diskown.disk events	718
diskown.diskhaswronginfo events	718
diskown.disknotsup events	719
diskown.diskwasstolen events	719
diskown.drawer events	719
diskown.duplicatesoftentries events	721
diskown.errorduringio events	721
diskown.errorreadingownership events	722
diskown.generated events	722
diskown.hlfshlf events	723
diskown.ilv events	724
diskown.incorrectownername events	725
diskown.isbeingenabled events	726
diskown.isenabled events	726
diskown.mcc events	727
diskown.mccip events	727
diskown.ownerchangefailed events	728
diskown.ownernameupdatefailed events	729
diskown.ownerreservationmismatch events	729
diskown.raid events	730

diskown.reassignfailed events	730
diskown.releasingmismatchedreservation events	731
diskown.replayingownership events	731
diskown.rescan events	732
diskown.rescanmessagefailed events	733
diskown.rescanwait events	733
diskown.rewriteowner events	734
diskown.shreddisk events	734
diskown.shlf events	735
diskown.split events	735
diskown.statusreadownership events	736
diskown.stealingdisk events	736
diskown.stealingdiskprevented events	737
diskown.uprevversion events	737
diskown.waitingforreassign events	738
diskshelf events	739
diskshelf.psu events	739
diskshelf.support events	739
diskshelf.supported events	740
diskshelf.unsupported events	740
dlock events	742
dlock.release events	742
dm events	743
dm.cluster events	743
dm.vserver events	743
dml events	745
dml.drive events	745
dml.fae events	745
dml.init events	746
dml.initialized events	746
dms events	747
dms.restserver events	747
dns events	748
dns.disabled events	748
dns.server events	748
download events	749
download.request events	749
download.requestdone events	749
download.requestdoneerror events	749
ds events	751
ds.sas events	751
dump events	757
dump.bitmap events	757
dump.fail events	757
dump.file events	757

dump.inconsistent events	758
dump.inconsistent_streamdir_lunclone events	760
dump.max events	760
dump.notify events	761
dump.register events	761
dump.snapshot events	762
dump.terminate events	763
dump.transit events	764
dump.volume events	764
dump.xfercomplete events	764
ekmip events	766
ekmip.akv events	766
ekmip.awskms events	769
ekmip.gckms events	771
ekmip.ikpkms events	774
ems events	778
ems.asup events	778
ems.engine events	778
ems.snmp events	780
ems.test events	780
enc events	782
enc.operation events	782
encryption events	785
encryption.op events	785
encryption.queue events	787
esh events	788
esh.async events	788
esh.auto events	788
esh.bypass events	789
esh.element events	790
esh.init events	791
esh.module events	791
esh.pctl events	791
esh.speed events	793
esh.term events	794
esh.unable events	794
esh.unexpected events	794
esh.xfer events	795
exports events	798
exports.anon events	798
exports.dns events	798
exports.dom events	799
exports.hostname events	800
exports.netgroup events	801
exports.ngbh events	802

exports.nsdbs events	803
exports.policy events	803
extcache events	805
extcache.aggregate events	805
extcache.configunsupported events	805
extcache.control events	805
extcache.disabled events	806
extcache.enabled events	806
extcache.io events	807
extcache.offline events	808
extcache.qos events	809
extcache.restart events	809
extcache.rw events	810
extcache.unallowed events	819
fabric events	821
fabric.mgmt events	821
fabric.sdm events	821
fabriclink events	822
fabriclink.retry events	822
fabricpool events	823
fabricpool.cloud events	823
fabricpool.full events	823
fabricpool.lost events	824
fabricpool.nearly events	824
fabricpool.scu events	825
fabricpool.tiering events	825
fal_nvme events	827
fal_nvme.init events	827
fal_nvme.online events	827
fal_nvme.partition events	827
fc events	829
fc.bridge events	829
fcadmin events	835
fcadmin.bridge events	835
fci events	839
fci.adapter events	839
fci.config events	846
fci.device events	848
fci.discovery events	850
fci.fabric events	851
fci.flash events	852
fci.fwupdate events	853
fci.link events	855
fci.loop events	856
fci.mserr events	856

fci.nserr events	857
fci.set events	859
fci.sfp events	859
fci.wrong events	860
fcmon events	861
fcmon.status events	861
fcpl events	862
fcpl.service events	862
fcvi events	863
fcvi.attr events	863
fcvi.cm events	863
fcvi.cross events	864
fcvi.disabled events	864
fcvi.fcrenabled events	864
fcvi.initfailed events	865
fcvi.invalid events	865
fcvi.link events	866
fcvi.qlqc events	867
fcvi.qlqc_regdumpfailed events	884
fcvi.unsupported events	885
fcvi.update events	885
ffm events	886
ffm.aggr events	886
ffm.attach events	887
ffm.dataip events	888
ffm.detach events	890
ffm.init events	891
ffm.restart events	891
ffrespmon events	893
ffrespmon.aggr events	893
ffrespmon.nfs events	893
fg events	895
fg.admin events	895
fg.aggr events	895
fg.convert events	896
fg.create events	898
fg.inactive events	899
fg.inodes events	900
fg.member events	901
fg.modify events	902
fg.move events	903
fg.multidelete events	903
fg.plex events	903
fg.provision events	904
fg.quota events	904

fg.recover events	906
fg.space events	907
fg.statechange events	909
fg.svmdr events	909
fg.unconfig events	910
fg.unconvert events	910
fg.writeback events	912
fgr events	913
fgr.metadata events	913
fips events	914
fips.post events	914
firmware events	915
firmware.file events	915
flexcache events	916
flexcache.attrib events	916
flexcache.autogrow events	917
flexcache.cacheconnected events	917
flexcache.cachedisconnected events	918
flexcache.capacity events	918
flexcache.conf events	919
flexcache.config events	919
flexcache.encryption events	920
flexcache.featuredisabled events	920
flexcache.mcc events	921
flexcache.name events	921
flexcache.nearly events	922
flexcache.origin events	922
flexcache.origincleanup events	924
flexcache.originconnected events	925
flexcache.origindisconnected events	926
flexcache.originfull events	926
flexcache.originfull_1 events	927
flexcache.softevict events	927
fli events	929
fli.lun events	929
fm events	932
fm.cpuutilduringarl events	932
fm.cpuutilduringtransition events	932
fm.overridebxbuncertain events	933
fmmb events	934
fmmb.add events	934
fmmb.blobnotfound events	934
fmmb.blobnotwritten events	934
fmmb.blobwritten events	935
fmmb.current events	936

fmmb.destroy events	936
fmmb.disk events	937
fmmb.exthdr events	937
fmmb.inststat events	938
fmmb.lock events	938
fmmb.major events	939
fmmb.mediator events	939
fmmb.mmr events	940
fmmb.newdisksuggestion events	941
fmmb.oldformatcleared events	941
fmmb.opaque events	942
fmmb.payload events	942
fmmb.revert events	942
fp events	944
fp.est events	944
fpolicy events	946
fpolicy.eagain events	946
fpolicy.fgconversion events	947
fpolicy.memory events	947
fpolicy.outage events	947
fpolicy.pspath events	951
fpolicy.server events	951
fsm events	955
fsm.low events	955
gb events	956
gb.cfo events	956
gb.netra events	962
gb.sfo events	963
geom events	973
geom.ontap events	973
gpo events	974
gpo.cap events	974
gpo.frmwork events	974
gpo.job events	976
gpo.reg events	978
gpo.sec events	979
gpo.security events	986
ha events	987
ha.aggrmigrduration events	987
ha.bdfu events	987
ha.disasterbootinitiated events	988
ha.fm events	988
ha.gb events	988
ha.gettakoverlockstatus events	990
ha.giveback events	991



ha.localnodeshutdown events	995
ha.mbxresvreleasedelay events	995
ha.netpartition events	996
ha.netparttodisabled events	997
ha.netra events	998
ha.nomediatormbx events	999
ha.noprogrammaticto events	999
ha.panicinfontsent events	999
ha.panicinfontsentic events	1000
ha.panicinforcvbadvrn events	1001
ha.panicinfosent events	1001
ha.panicinfosznotok events	1002
ha.resvconflicthalt events	1002
ha.resvmediatormbxstatus events	1003
ha.takeover events	1003
ha.takeoverimpdegraded events	1005
ha.takeoverimphotshelf events	1005
ha.takeoverimpic events	1005
ha.takeoverimplowmem events	1006
ha.takeoverimplotdef events	1006
ha.takeoverimpunsync events	1007
ha.takeoverimpversion events	1007
ha.tkover events	1008
ha.tkoverabortednomediatormbx events	1009
ha.toafterrcvpanicinfo events	1009
ha.tp events	1009
hamsg events	1011
hamsg.diskcheckresp events	1011
hamsg.dskchkpopulateerr events	1011
haosc events	1012
haosc.config events	1012
haosc.fru events	1013
haosc.invalid events	1013
haosc.sa events	1014
hashd events	1015
hashd.maxstore events	1015
hm events	1016
hm.alert events	1016
host events	1018
host.app events	1018
host.error events	1019
hotplug events	1020
hotplug.generic events	1020
hotplug.insert events	1020
hotplug.remove events	1022

hotplug.replace events	1023
http events	1024
http.cfg events	1024
httpd events	1025
httpd.api events	1025
hya events	1026
hya.dup_vbn_lookup events	1026
hybaggrcopy events	1027
hybaggrcopy.dstgeomismatch events	1027
hybaggrcopy.srcgeomismatch events	1027
ic events	1028
ic.hainterconnectdown events	1028
ic.hainterconnectlinkdown events	1028
ic.linkspeeddegraded events	1029
ic.linkstatuschange events	1029
ic.linkwidthdegraded events	1030
ic.unsupportcable events	1030
ic.viif events	1031
iclink events	1032
iclink.init events	1032
iclink.probe events	1032
icmp events	1033
icmp.ping events	1033
image events	1034
image.install events	1034
ipfw events	1036
ipfw.reachedmaxstates events	1036
ipmultipath events	1037
ipmultipath.init events	1037
ipmultipath.session events	1037
ipmultipath.setup events	1038
ipsec events	1040
ipsec.ib events	1040
ipsec.ob events	1041
iscsi events	1043
iscsi.loginfailure events	1043
iscsi.loginipnotauthorized events	1043
iscsi.notice events	1044
iscsi.session events	1044
iscsi.warning events	1044
isns events	1046
isns.enqueue events	1046
isns.portal events	1046
isns.reg events	1046
isns.socket events	1048

kern events	1050
kern.cron events	1050
kern.log events	1050
kern.rc events	1051
kern.shutdown events	1052
kern.syslog events	1052
kern.syslogd events	1053
kern.timezone events	1054
kern.uptime events	1054
kern.version events	1055
kern.vm events	1055
km events	1056
km.cluster events	1056
km.cmek events	1056
km.extkeysvr events	1058
km.keyretrieve events	1059
km.keyserver events	1059
km.mcc events	1060
km.okmdb events	1062
km.onboard events	1062
km.restore events	1063
km.run events	1063
km.volume events	1066
kmip2 events	1068
kmip2.cannot events	1068
kmip2.invalid events	1069
kmip2.ssl events	1070
krb events	1071
krb.config events	1071
ktls events	1072
ktls.badauth events	1072
ktls.cnxnhandshakelimit events	1072
ktls.failed events	1073
ktls.health events	1073
largevol events	1079
largevol.remount events	1079
ldap events	1080
ldap.false events	1080
lib events	1081
lib.nvtrace events	1081
lic events	1082
lic.cappool events	1082
lic.lm events	1084
license events	1085
license.capac events	1085

license.cappool events	1091
license.cert events	1092
license.check events	1093
license.feat events	1096
license.fpool events	1098
license.ontapi events	1100
license.pool events	1101
license.sfo events	1101
license.subscription events	1102
license.term events	1103
lifp events	1105
lifp.net events	1105
lldp events	1106
lldp.receive events	1106
lmgr events	1107
lmgr.aggr events	1107
lmgr.badstructure events	1107
lmgr.ca events	1108
lmgr.cb events	1108
lmgr.delegations events	1109
lmgr.files events	1110
lmgr.gb events	1110
lmgr.hostfile events	1111
lmgr.hosts events	1112
lmgr.lock events	1113
lmgr.locks events	1114
lmgr.lockstate events	1116
lmgr.mcc events	1117
lmgr.meta events	1117
lmgr.mismatch events	1118
lmgr.move events	1118
lmgr.notify events	1119
lmgr.owners events	1119
lmgr.partner events	1120
lmgr.persistent events	1122
lmgr.precommit events	1122
lmgr.pruning events	1123
lmgr.recons events	1124
lmgr.sendhome events	1124
lmgr.session events	1125
lmgr.sf events	1125
lmgr.sh events	1126
lmgr.vol events	1126
logger events	1128
logger.usr events	1128

login events	1129
login.auth events	1129
lsm events	1130
lsm.merge events	1130
lsm.pal events	1131
lsm.rangedelete events	1131
lun events	1133
lun.clone events	1133
lun.copy events	1134
lun.create events	1134
lun.data events	1135
lun.destroy events	1136
lun.inaccessible events	1136
lun.inconsistent events	1137
lun.inplace events	1138
lun.metadb events	1140
lun.metafile events	1140
lun.move events	1141
lun.no events	1142
lun.nvfail events	1143
lun.offline events	1144
lun.online events	1145
lun.oovc events	1145
lun.op events	1146
lun.out events	1147
lun.outplace events	1147
lun.provisioning events	1148
lun.restore events	1148
lun.scsi events	1149
lun.serial events	1149
lun.space events	1150
lun.transition events	1150
lun.vdisk events	1154
lun.vol events	1154
lun.vtoc events	1155
malloc events	1157
malloc.failure events	1157
master events	1158
master.cast events	1158
master.couldnt events	1158
master.demote events	1159
master.new events	1160
master.releasing events	1161
master.taking events	1161
master.try events	1161

maxcid events	1163
maxcid.limit events	1163
maxcid.threshold events	1163
mc events	1164
mc.channel events	1164
mc.post events	1164
mcc events	1165
mcc.auso events	1165
mcc.check events	1166
mcc.config events	1166
mcc.ctlbe events	1176
mcc.drdom events	1177
mcc.drgroup events	1178
mcc.drmsg events	1179
mcc.drsom events	1180
mcc.fabricpool events	1185
mcc.forced events	1187
mcc.hwassist events	1187
mcc.ic events	1188
mcc.ipspace events	1188
mcc.kill events	1189
mcc.lunmgr events	1189
mcc.network events	1189
mcc.notify events	1190
mcc.nso events	1190
mcc.nvmeof events	1191
mcc.sb events	1193
mcc.sbca events	1194
mcc.scsipd events	1195
mcc.so events	1195
mcc.switchover events	1195
mcc.toomany events	1196
mct events	1197
mct.channel events	1197
mct.control events	1197
mdb events	1198
mdb.exit events	1198
mes_dimm events	1199
mes_dimm.error events	1199
mgmt events	1200
mgmt.aba events	1200
mgmt.abv events	1203
mgmt.cifsserver events	1205
mgmt.cifsshare events	1206
mgmt.cifssuperuser events	1207

mgmt.cifssymlink events	1208
mgmt.cr events	1209
mgmt.exovol events	1212
mgmt.fpolicy events	1213
mgmt.fsecurity events	1215
mgmt.keymgr events	1215
mgmt.mcc events	1216
mgmt.nfslisteners events	1216
mgmt.nfsserver events	1217
mgmt.rst events	1218
mgmt.s3 events	1218
mgmt.sm events	1219
mgmt.smbc events	1230
mgmt.snapmir events	1231
mgmt.vdr events	1237
mgmt.vopl events	1238
mgmtgwd events	1265
mgmtgwd.certificate events	1265
mgmtgwd.configbr events	1266
mgmtgwd.fg events	1270
mgmtgwd.filereplication events	1273
mgmtgwd.fips events	1274
mgmtgwd.jobmgr events	1275
mgmtgwd.mixed events	1279
mgmtgwd.replace events	1280
mgmtgwd.repos events	1280
mgmtgwd.revertto events	1281
mgmtgwd.rootvol events	1281
mgmtgwd.rootvolrec events	1283
mgmtgwd.schedsnap events	1284
mgmtgwd.sign events	1285
mgmtgwd.snapshot events	1286
mgmtgwd.upgrade events	1288
mgmtgwd.var events	1289
mgmtgwd.volcreatedrestricted events	1289
mgmtgwd.volmove events	1290
mgmtgwd.vreport events	1291
mgmtgwd.wafl events	1291
mgr events	1293
mgr.avail events	1293
mgr.boot events	1293
mgr.module events	1299
mgr.partner events	1300
mgr.stack events	1300
mgwd events	1305

mgwd.cluscheck events	1305
mgwd.notify events	1305
mhost events	1307
mhost.ca events	1307
migrate events	1310
migrate.root events	1310
misconfig events	1312
misconfig.lif events	1312
mlm events	1313
mlm.arr events	1313
mlm.array events	1313
mlm.arraylun events	1334
mlm.bad events	1335
mlm.excessive events	1335
mlm.non events	1336
mlm.prreregistration events	1336
mlm.scsi events	1337
mlm.tapemc events	1338
mlm.tpqd events	1341
mlm.vmdisk events	1342
monitor events	1343
monitor.brokendisk events	1343
monitor.chassisfan events	1343
monitor.chassisfanfail events	1345
monitor.chassispower events	1345
monitor.chassispowersupplies events	1346
monitor.chassispowersupply events	1346
monitor.chassistemperature events	1348
monitor.completed events	1349
monitor.cpu events	1350
monitor.disklabelcheckfailed events	1351
monitor.downrevdisks events	1351
monitor.extcache events	1352
monitor.fan events	1352
monitor.fru events	1354
monitor.globalstatus events	1354
monitor.iocard events	1356
monitor.ioexpansionpower events	1357
monitor.ioexpansiontemperature events	1357
monitor.ioxmtemp events	1359
monitor.mismatch events	1359
monitor.nvmembattery events	1360
monitor.nvramlowbatteries events	1360
monitor.nvramlowbattery events	1361
monitor.partnercontroller events	1362



monitor.power events	1362
monitor.raid events	1364
monitor.raiddp events	1365
monitor.raidtec events	1366
monitor.rtchighpower events	1366
monitor.rtclopower events	1366
monitor.rtcnormal events	1367
monitor.rtcwarnlowpower events	1367
monitor.shelf events	1368
monitor.shutdown events	1371
monitor.sparelabelcheckfailed events	1376
monitor.temp events	1377
monitor.unknown events	1377
monitor.vol events	1378
monitor.volume events	1379
monitor.weeklystatsasup events	1380
moveagent events	1382
moveagent.abort events	1382
moveagent.cleanup events	1382
moveagent.commitdst events	1383
moveagent.commitsrc events	1383
moveagent.cutovercomp events	1384
moveagent.detected events	1384
moveagent.dst events	1385
moveagent.exceed events	1385
moveagent.outofmemory events	1386
moveagent.phase events	1386
moveagent.swapvldb events	1387
moveagent.switchback events	1388
moveagent.switchover events	1388
mpt events	1390
mpt.action events	1390
mpt.bus events	1390
mpt.busreg events	1391
mpt.path events	1391
mpt.phydisk events	1392
mpt.tmf events	1392
msata events	1393
msata.boot events	1393
msata.unable events	1393
msrespmon events	1395
msrespmon.nfs events	1395
multi events	1396
multi.dsk events	1396
mw events	1397

mw.hir events	1397
mw.iron events	1397
nameserv events	1398
nameserv.conf events	1398
nameserv.config events	1398
nameserv.ddns events	1399
nameserv.netgrpbyhost events	1401
nand events	1403
nand.predict events	1403
nand.unable events	1403
nblade events	1404
nblade.aclcompliance events	1404
nblade.callbacktimedout events	1404
nblade.cifs events	1405
nblade.cifscsaccessdenied events	1406
nblade.cifscsaccessdenied events	1406
nblade.cifscsdactionfailed events	1407
nblade.cifscsdmaxtxnlimit events	1407
nblade.cifslongrunpattmatch events	1407
nblade.cifsmanyauths events	1408
nblade.cifsmaxopensamefile events	1409
nblade.cifsmaxopensamefilenotice events	1409
nblade.cifsmaxsessperusrconn events	1410
nblade.cifsmaxsessperusrconnnotice events	1410
nblade.cifsmaxwatchespertree events	1411
nblade.cifsmaxwatchespertreenotice events	1412
nblade.cifsmemexceeded events	1412
nblade.cifsnbnameconflict events	1413
nblade.cifsnonaesnibytesproc events	1413
nblade.cifsnonunicoderequest events	1414
nblade.cifsnoprivshare events	1414
nblade.cifsoperationtimedout events	1415
nblade.cifssametreepersess events	1415
nblade.cifssametreepersessnotice events	1416
nblade.cifsshconnectfailed events	1417
nblade.cifswitnessfonotify events	1417
nblade.css events	1418
nblade.cvo events	1418
nblade.dbladenoreponse events	1419
nblade.didnotinitialize events	1420
nblade.ecv events	1420
nblade.execoverlimit events	1421
nblade.exportaccesschkfailed events	1421
nblade.exportaccessindeterm events	1422
nblade.fcvoldisconnected events	1422

nblade.fileopenlimitexceeded events . . . . .	1423
nblade.flexcachecaaccess events . . . . .	1423
nblade.flexcachevolumeaccess events . . . . .	1424
nblade.flexgroupstatefulprotocolaccess events . . . . .	1424
nblade.fpolicy events . . . . .	1425
nblade.fpolpassthruconnect events . . . . .	1426
nblade.fpolpassthrudisconnect events . . . . .	1427
nblade.fpolpassthruonsmb1 events . . . . .	1427
nblade.fpolpsinitfail events . . . . .	1427
nblade.fpolpsvolmoving events . . . . .	1428
nblade.fpolpsvolnoaccess events . . . . .	1428
nblade.fpolpsvolnospace events . . . . .	1429
nblade.fpolpsvolnotfound events . . . . .	1429
nblade.fpolpsvolofflined events . . . . .	1430
nblade.gracebegin events . . . . .	1430
nblade.graceend events . . . . .	1431
nblade.httpmaxconntimeouts events . . . . .	1431
nblade.httpmemexceeded events . . . . .	1432
nblade.invtcprecordmarker events . . . . .	1432
nblade.junctionrootlookup2 events . . . . .	1433
nblade.longrunningrpcop events . . . . .	1433
nblade.nbnstoomanylifs events . . . . .	1434
nblade.newclientidmismatch events . . . . .	1434
nblade.nfs4illegaldirentname events . . . . .	1435
nblade.nfs4opnotsupported events . . . . .	1435
nblade.nfs4sequenceinvalid events . . . . .	1436
nblade.nfs4trunkmaxlimit events . . . . .	1436
nblade.nfs4unoptimizedio events . . . . .	1437
nblade.nfsconnresetandclose events . . . . .	1437
nblade.nfscredcacheflushed events . . . . .	1438
nblade.nfsmountrootonly events . . . . .	1439
nblade.nfspathresmaxlinks events . . . . .	1439
nblade.nfsrcchecksummismatch events . . . . .	1440
nblade.nfsv3writetoolarge events . . . . .	1440
nblade.nfsv4blockclient events . . . . .	1441
nblade.nfsv4nsbdbdomainmismatch events . . . . .	1441
nblade.nfsv4pool exhaust events . . . . .	1442
nblade.nfsv4poolthreshold events . . . . .	1442
nblade.nfsv4writetoolarge events . . . . .	1443
nblade.nocsmssession events . . . . .	1443
nblade.nosmbvernegotiated events . . . . .	1444
nblade.notactivevscanserver events . . . . .	1444
nblade.pcp events . . . . .	1445
nblade.qosvioldetectregfail events . . . . .	1445
nblade.rcbinid events . . . . .	1446

nblade.recoverybegin events . . . . .	1446
nblade.s3bucketpathnotfound events . . . . .	1446
nblade.scannerconnected events . . . . .	1447
nblade.scannerdisconnected events . . . . .	1447
nblade.sectracequeueoverload events . . . . .	1448
nblade.sharemaxconnlimit events . . . . .	1448
nblade.showmountdisabled events . . . . .	1449
nblade.smb1onlyclientfound events . . . . .	1449
nblade.smbsignaturemismatch events . . . . .	1449
nblade.testevent events . . . . .	1450
nblade.vifdel events . . . . .	1450
nblade.vldb events . . . . .	1451
nblade.vscanbadipprivaccess events . . . . .	1452
nblade.vscanbadprotomagicnum events . . . . .	1453
nblade.vscanbaduserprivaccess events . . . . .	1453
nblade.vscanconnbackpressure events . . . . .	1454
nblade.vscanconninactive events . . . . .	1454
nblade.vscanconninvaliduser events . . . . .	1455
nblade.vscanconnreqonsmb1 events . . . . .	1455
nblade.vscanexcessivetos events . . . . .	1456
nblade.vscannodispatcher events . . . . .	1456
nblade.vscannopolicyenabled events . . . . .	1457
nblade.vscannoregdscanner events . . . . .	1457
nblade.vscannoscannerconn events . . . . .	1458
nblade.vscanvirusdetected events . . . . .	1458
nblade.vscanworkqueueoverloaded events . . . . .	1459
ndmp events . . . . .	1460
ndmp.ctrl events . . . . .	1460
ndmp.data events . . . . .	1460
ndmp.fail events . . . . .	1461
ndmp.max events . . . . .	1461
ndmp.password events . . . . .	1461
ndo events . . . . .	1463
ndo.cr events . . . . .	1463
net events . . . . .	1474
net.cg events . . . . .	1474
net.fib events . . . . .	1475
net.ifconfig events . . . . .	1476
net.ifgrp events . . . . .	1477
net.ipconfig events . . . . .	1479
net.static events . . . . .	1479
net.vserverupdatefailure events . . . . .	1480
net.vserverupdatesuccess events . . . . .	1480
netgroup events . . . . .	1481
netgroup.files events . . . . .	1481

netgroup.ldap events	1481
netgroup.nis events	1482
netif events	1484
netif.autonegfailed events	1484
netif.badeeprom events	1484
netif.badoversizepacket events	1484
netif.cprevented events	1485
netif.cprevented detected events	1485
netif.cprmaxeventsreached events	1486
netif.eepromreadfailed events	1486
netif.excesslinkinterrupts events	1487
netif.fatal events	1487
netif.fwupdatefailed events	1487
netif.fwupdatestatus events	1488
netif.hangdetected events	1488
netif.hn events	1489
netif.init events	1489
netif.linkdown events	1490
netif.linkerrors events	1490
netif.linkinfo events	1491
netif.linkup events	1491
netif.nicnotsupported events	1492
netif.overtemperror events	1492
netif.phynotread events	1492
netif.phyoperationfailed events	1493
netif.phyread events	1493
netif.ratelimitthreshold events	1494
netif.sfpeventerrorcode events	1494
netif.sfpnotsupported events	1495
netif.speeddowngraded events	1495
netif.speednotsupported events	1496
netif.tcp events	1496
netif.toomanynics events	1496
netif.uncoreccerror events	1497
netif.unknownswrequest events	1497
netinet events	1499
netinet.ethr events	1499
netinet.icmp events	1500
netinet6 events	1502
netinet6.frag events	1502
netinet6.icmp events	1502
netinet6.in6 events	1504
netinet6.input events	1505
netinet6.mld6 events	1506
netinet6.nbr events	1507

netinet6.nd6 events	1514
netinet6.rtr events	1515
netport events	1519
netport.eth events	1519
netport.pvlan events	1519
nfs events	1520
nfs.client events	1520
nfs.krb events	1520
nis events	1521
nis.active events	1521
nis.db events	1521
nis.group events	1521
nis.netgrp events	1522
nis.server events	1524
no events	1526
no.halt events	1526
noderun events	1529
noderun.login events	1529
nodewatchdog events	1530
nodewatchdog.config events	1530
nodewatchdog.gfc events	1530
nodewatchdog.node events	1531
nodewatchdog.resrc events	1532
nodewatchdog.svc events	1533
nonha events	1535
nonha.resvconflicthalt events	1535
nse events	1536
nse.authlockout events	1536
nse.op events	1536
nv events	1538
nv.data events	1538
nv.extension events	1538
nv.fake events	1539
nv.fio events	1540
nv.flash events	1540
nv.none events	1541
nv.partner events	1542
nv.rename events	1543
nv.switchover events	1544
nvd events	1546
nvd.biodone events	1546
nvdimm events	1547
nvdimm.excessive events	1547
nvdimm.flash events	1547
nvdimm.fw events	1547

nvdimm.nvemmem events	1548
nvdimm.nvmem events	1549
nvlog events	1552
nvlog.replay events	1552
nvme events	1553
nvme.boot events	1553
nvme.card events	1553
nvme.fw events	1554
nvme.link events	1555
nvme.ssd events	1556
nvmem events	1558
nvmem.battery events	1558
nvmem.voltage events	1577
nvmens events	1580
nvmens.512blksz events	1580
nvmens.destroy events	1580
nvmens.offline events	1581
nvmens.online events	1581
nvmens.out events	1581
nvmens.repl events	1582
nvmens.smbc events	1582
nvmeof events	1584
nvmeof.cabling events	1584
nvmeof.hostport events	1584
nvmeof.mixedports events	1585
nvmeof.onlyoneport events	1586
nvmeof.port events	1587
nvmeof.portioma events	1587
nvmeof.portiomb events	1588
nvmeof.rdma events	1589
nvmeof.subsystem events	1590
nvmeof.timeout events	1590
nvmf events	1592
nvmf.graceperiod events	1592
nvmf.mgmt events	1593
nvmf.qos events	1595
nvmf.remote events	1596
nvmf.roce events	1596
nvmf.tcp events	1597
nvpm events	1598
nvpm.extend events	1598
nvrmm events	1599
nvrmm.battery events	1599
nvrmm.excessive events	1616
nvrmm.flash events	1617

nvrnm.hw events	1618
nwd events	1620
nwd.dataip events	1620
oauth2 events	1621
oauth2.httppd events	1621
oauth2.localroles events	1621
oauth2.status events	1622
oauth2.wsdm events	1622
object events	1624
object.store events	1624
objectattach events	1631
objectattach.hw events	1631
objstore events	1632
objstore.host events	1632
objstore.interclusterlifedown events	1632
old events	1633
old.replica events	1633
openssh events	1634
openssh.unsupported events	1634
osc events	1636
osc.bucket events	1636
osc.signaturremismatch events	1636
other events	1637
other.veto events	1637
pac events	1639
pac.access events	1639
pac.if events	1639
pac.peer events	1640
pac.sock events	1640
passwd events	1641
passwd.changed events	1641
pci events	1642
pci.init events	1642
pcie events	1643
pcie.errors events	1643
pcie.link events	1643
pcnfsd events	1644
pcnfsd.auth events	1644
pcp events	1645
pcp.linkstatechangesucceed events	1645
pcp.linkstateinprogress events	1645
perclientstorepoolthreshold events	1646
perclientstorepoolthreshold events	1646
perf events	1647
perf.ccma events	1647



perfstatd events	1649
perfstatd.dir events	1649
perfstatd.disk events	1649
perfstatd.start events	1650
pha events	1651
pha.add events	1651
pha.dev events	1651
pha.disk events	1654
pha.obj events	1654
pha.vsa events	1659
phf events	1660
phf.manage events	1660
platform events	1661
platform.chsismobonamediff events	1661
platform.dimmconfigerror events	1661
platform.insufficientmemory events	1661
platform.nopartnerpresent events	1662
platform.reducedmemory events	1662
pn events	1664
pn.driver events	1664
pn.mvia events	1664
probedone events	1666
probedone.inquiry events	1666
protocol events	1667
protocol.exportcheckfailed events	1667
prov events	1668
prov.rollback events	1668
pubsub events	1670
pubsub.kernel events	1670
pvif events	1672
pvif.badmtu events	1672
pvif.lacp events	1672
qos events	1673
qos.bgtask events	1673
qos.mgmt events	1673
qos.min events	1674
qos.mintp events	1674
qos.monitor events	1675
qos.onnix events	1676
qos.qtree events	1676
qos.viodet events	1676
qos.violation events	1678
qtree events	1680
qtree.delete events	1680
query events	1682

query.execution events	1682
quota events	1683
quota.inconsistent events	1683
quota.intr events	1683
quota.parse events	1684
quota.push events	1684
quota.reinit events	1685
quota.resize events	1685
quota.scan events	1686
quota.softlimit events	1688
quota.upgrade events	1689
raid events	1690
raid.aggr events	1690
raid.aggregate events	1691
raid.aggrvote events	1691
raid.assim events	1693
raid.autopart events	1710
raid.capacity events	1713
raid.carrier events	1714
raid.cksum events	1716
raid.config events	1747
raid.data events	1774
raid.debug events	1777
raid.disable events	1777
raid.disk events	1778
raid.diskadd events	1806
raid.disktoc events	1808
raid.enable events	1809
raid.encrypt events	1810
raid.fdr events	1810
raid.flexvol events	1813
raid.fm events	1814
raid.foreignaggregate events	1818
raid.fs events	1819
raid.hybrid events	1819
raid.initialize events	1820
raid.label events	1821
raid.labeledit events	1824
raid.labelmaint events	1824
raid.mcc events	1826
raid.mgrio events	1831
raid.mgvio events	1832
raid.mirror events	1832
raid.monitor events	1854
raid.multierr events	1855

raid.mv events	1862
raid.no events	1864
raid.ns events	1864
raid.nvram events	1865
raid.offline events	1868
raid.olm events	1868
raid.partition events	1869
raid.plex events	1870
raid.pmdopt events	1871
raid.poolsort events	1873
raid.preserve events	1875
raid.quiesce events	1876
raid.read events	1876
raid.readerr events	1880
raid.recons events	1886
raid.reparity events	1887
raid.replay events	1887
raid.rg events	1889
raid.root events	1950
raid.scrub events	1951
raid.shared events	1953
raid.sp events	1955
raid.sparecore events	1960
raid.spares events	1960
raid.stripe events	1964
raid.switchoverfail events	1969
raid.switchovernodisks events	1970
raid.syncmirr events	1970
raid.syncmirror events	1970
raid.tetris events	1972
raid.tree events	1978
raid.uninitialized events	1978
raid.unpartition events	1979
raid.unsupported events	1980
raid.verify events	1981
raid.vldb events	1981
raid.vol events	1982
raid.warn events	1996
raidlm events	1997
raidlm.cannotmakeprogress events	1997
raidlm.carrier events	1997
rbac events	2001
rbac.spuser events	2001
rdb events	2002
rdb.corruption events	2002

rdb.env events	2002
rdb.ha events	2004
rdb.internal events	2005
rdb.node events	2005
rdb.open events	2006
rdb.recovery events	2006
rdb.sql events	2007
rdma events	2009
rdma.rlib events	2009
rdmalib events	2011
rdmalib.abort events	2011
rdmalib.duplicateid events	2011
rdmalib.multichan events	2011
rdmalib.rcm events	2012
recover events	2013
recover.abort events	2013
reg events	2014
reg.cookie events	2014
reg.errno events	2014
reg.file events	2015
reg.key events	2018
reg.options events	2019
reg.transaction events	2020
repl events	2023
repl.archival events	2023
repl.autol2l events	2023
repl.cg events	2023
repl.checker events	2024
repl.ckpt events	2025
repl.cloud events	2025
repl.core events	2026
repl.data events	2027
repl.dest events	2027
repl.dst events	2029
repl.enable events	2030
repl.exovol events	2030
repl.extended events	2031
repl.gsr events	2031
repl.holdblock events	2032
repl.hole events	2034
repl.jumpahead events	2034
repl.large events	2035
repl.logical events	2037
repl.out events	2038
repl.physdiff events	2039

repl.s2c events	2040
repl.send events	2043
repl.seq events	2044
repl.ser events	2044
repl.space events	2045
repl.src events	2045
repl.switch events	2046
repl.unable events	2046
repl.xfer events	2047
replaycache events	2048
replaycache.arcbkallocerr events	2048
replaycache.nfsv41bindelete events	2048
replication events	2050
replication.cluster events	2050
replication.dst events	2050
replication.err events	2057
replication.givebackcancel events	2058
replication.global events	2058
replication.lcl events	2059
replication.nomemory events	2059
replication.qtree events	2060
replication.regoperr events	2062
replication.resync events	2062
replication.rmt events	2063
replication.src events	2063
replication.stats events	2066
replication.status events	2066
replication.upgrade events	2067
replication.vol events	2067
replication.wormlog events	2070
replication.xfr events	2070
repository events	2072
repository.cleanup events	2072
repository.create events	2072
repository.ctpupdate events	2073
repository.dangling events	2073
repository.dc events	2074
repository.inconsistent events	2075
repository.invalid events	2075
repository.ns events	2076
repository.snapshot events	2077
respmon events	2079
respmon.service events	2079
revert events	2080
revert.ufmt events	2080

revertto events	2084
revertto.force events	2084
revertto.start events	2084
revertto.system events	2084
rlib events	2086
rlib.duplicate events	2086
rlib.ifconfig events	2086
rlib.mccip events	2086
rlib.remote events	2087
rlm events	2089
rlm.driver events	2089
rlm.firmware events	2090
rlm.heartbeat events	2091
rlm.network events	2092
rlm.notconfigured events	2093
rlm.orftp events	2093
rlm.snmp events	2094
rlm.ssh events	2094
rlm.userlist events	2094
rpc events	2096
rpc.unspecified events	2096
rpc.xprt events	2096
rpl events	2098
rpl.check events	2098
rpl.fixer events	2098
rpl.snapshot events	2098
rpl.src events	2099
rpt events	2100
rpt.disk events	2100
rsh events	2101
rsh.enabled events	2101
rsh.firewall events	2101
rsh.rcmd events	2101
rsh.setsockopt events	2102
rsh.socket events	2103
rshd events	2104
rshd.failure events	2104
s3 events	2105
s3.bucket events	2105
saml events	2107
saml.cfg events	2107
san events	2109
san.changedhomenode events	2109
san.changedhomeport events	2109
san.duplicateipaddress events	2109

san.duplicatewwpn events	2110
san.fcadapter events	2111
san.fcnvmeunsupportedport events	2111
san.initiator events	2112
san.iscsilifrepaired events	2113
san.lifmove events	2113
san.lifmoveafterswitchback events	2114
san.lifrepfieldsmissing events	2115
san.lifspernodelimit events	2115
san.mccpurgeverifyfailed events	2116
san.mccvserverinitfailed events	2116
san.mccvserverinitpending events	2117
san.missinglicense events	2117
san.noconnectedfcport events	2117
san.noconnecteddiscsiport events	2119
san.nofcportsonnode events	2119
san.nofcportwithadaptertype events	2120
san.noifgrpport events	2121
san.nonetworkportadaptertype events	2121
san.noportsinipospace events	2122
san.nosamenamedfcport events	2122
san.nosamenameddiscsiport events	2123
san.nosupportedssamenamedfcport events	2124
san.novlanport events	2124
san.optimizedpersonality events	2125
san.slmsupportcheckfailed events	2125
san.volrehostigroupcreatefailed events	2126
san.volrehostlunmapcreated events	2126
san.volrehostlunmapfailed events	2127
san.volrehostsanbestpractice events	2127
sas events	2129
sas.8001 events	2129
sas.8072 events	2129
sas.adapter events	2130
sas.cable events	2135
sas.config events	2136
sas.cpr events	2138
sas.device events	2139
sas.initialization events	2140
sas.link events	2141
sas.reset events	2141
sas.shelf events	2143
sas.unsupported events	2143
sasmon events	2145
sasmon.adapter events	2145

sasmon.disable events	2145
savecore events	2146
savecore.completed events	2146
savecore.nospace events	2146
savecore.readonly events	2146
savecore.waiting events	2147
sblade events	2148
sblade.duplicate events	2148
scsi events	2149
scsi.cmd events	2149
scsi.mcc events	2170
scsi.path events	2171
scsi.scan events	2172
scsi.security events	2172
scsi.sim events	2173
scsiblade events	2175
scsiblade.autoheal events	2175
scsiblade.cache events	2175
scsiblade.disabled events	2175
scsiblade.dpa events	2176
scsiblade.ha events	2176
scsiblade.illegal events	2177
scsiblade.import events	2177
scsiblade.init events	2179
scsiblade.invalid events	2180
scsiblade.isolated events	2180
scsiblade.kernel events	2181
scsiblade.lif events	2182
scsiblade.lu events	2183
scsiblade.lun events	2187
scsiblade.mapped events	2189
scsiblade.mcc events	2190
scsiblade.mgmt events	2190
scsiblade.object events	2192
scsiblade.offline events	2193
scsiblade.online events	2193
scsiblade.pport events	2193
scsiblade.prop events	2195
scsiblade.qosvioldetectregfail events	2195
scsiblade.reload events	2196
scsiblade.san events	2196
scsiblade.tiered events	2198
scsiblade.unavailable events	2198
scsiblade.vol events	2199
scsiblade.volume events	2199



scsiblade.vs events	2200
scsiblade.vserver events	2200
scsitarget events	2202
scsitarget.fcp events	2202
scsitarget.fct events	2203
scsitarget.hwpfct events	2210
scsitarget.ispfct events	2214
scsitarget.slifct events	2217
secd events	2222
secd.asidtrans events	2222
secd.authsys events	2222
secd.cifs events	2222
secd.cifsauth events	2223
secd.conn events	2225
secd.dc events	2225
secd.dns events	2226
secd.ipq events	2227
secd.kerberos events	2227
secd.ldap events	2230
secd.lsa events	2234
secd.nametrans events	2235
secd.netgroup events	2237
secd.netlogon events	2238
secd.nfs events	2238
secd.nfsauth events	2239
secd.nis events	2241
secd.nonetgroupfile events	2243
secd.quark events	2243
secd.rpc events	2244
secd.single events	2245
secd.strong events	2246
secd.unexpectedfailure events	2246
secd.unixlookupfailure events	2247
sectrace events	2248
sectrace.warning events	2248
securepurge events	2249
securepurge.aborted events	2249
securepurge.completed events	2249
securepurge.started events	2249
security events	2251
security.invalid events	2251
service events	2252
service.resp events	2252
service.unresp events	2252
ses events	2254

ses.access events	2254
ses.badsharestorageconfigerr events	2255
ses.bridge events	2256
ses.channel events	2257
ses.config events	2257
ses.disk events	2259
ses.download events	2260
ses.drive events	2261
ses.fw events	2263
ses.giveback events	2264
ses.inconsistent events	2265
ses.loop events	2265
ses.mismatch events	2267
ses.multipath events	2267
ses.psu events	2268
ses.shelf events	2269
ses.status events	2282
ses.system events	2310
ses.unsupported events	2311
sfo events	2313
sfo.aggr events	2313
sfo.arl events	2314
sfo.giveback events	2315
sfo.ignoretakeover events	2319
sfo.partialgiveback events	2320
sfo.reassignfailed events	2320
sfo.reassignnopartner events	2321
sfo.retry events	2321
sfo.retrying events	2322
sfo.sendhome events	2322
sfo.takenover events	2323
sfo.takeover events	2325
sfo.tkabort events	2329
sfo.tkovertoclusterwait events	2329
sfu events	2331
sfu.adaptersuspendio events	2331
sfu.auto events	2331
sfu.badbridgeshelfconfig events	2332
sfu.ctrllerlmntspershelf events	2333
sfu.download events	2333
sfu.downloadctrllerbridge events	2334
sfu.downloadererror events	2335
sfu.downloadingcontroller events	2335
sfu.downloadingctrlerr1xx events	2336
sfu.downloadstarted events	2336

sfu.downloadsucces events	2336
sfu.downloadsummary events	2337
sfu.downloadsummaryerrors events	2337
sfu.fcdownloadfailed events	2338
sfu.firmwaredownrev events	2338
sfu.firmwareuptodate events	2339
sfu.nsmdownloadfailed events	2340
sfu.partnerinaccessible events	2340
sfu.partnernotresponding events	2341
sfu.partnerrefusedupdate events	2341
sfu.partnerupdatecomplete events	2342
sfu.partnerupdatetimetype events	2342
sfu.rebootrequest events	2343
sfu.rebootrequestfailure events	2343
sfu.resumediskio events	2343
sfu.sasdownloadfailed events	2344
sfu.statuscheckfailure events	2344
sfu.suspenddiskio events	2345
sfu.suspendses events	2345
shelf events	2347
shelf.config events	2347
shelf.identify events	2354
shelf.io events	2355
shelf.module events	2355
shelf.reset events	2356
shelf.test events	2356
shm events	2358
shm.bypassed events	2358
shm.disk events	2359
shm.fab events	2359
shm.failurebytes events	2360
shm.nvme events	2360
shm.pulllogwarning events	2361
shm.reformat events	2361
shm.ssd events	2362
shm.threshold events	2363
showmount events	2371
showmount.update events	2371
sis events	2372
sis.aggr events	2372
sis.auto events	2373
sis.autosched events	2373
sis.c2c events	2373
sis.cfg events	2374
sis.changelog events	2375

sis.chkpoint events	2377
sis.clog events	2378
sis.compress events	2378
sis.corrupt events	2379
sis.default events	2379
sis.enable events	2380
sis.fingerprint events	2381
sis.idedup events	2381
sis.incorrect events	2382
sis.logical events	2382
sis.max events	2383
sis.metadata events	2384
sis.op events	2384
sis.policy events	2385
sis.reg events	2386
sis.revert events	2387
sis.revertmetacreate events	2389
sis.schedule events	2389
sis.undo events	2391
sis.upgrade events	2393
sis.upgrademetatile events	2395
sis.verify events	2396
sis.vol events	2396
sis.wafliron events	2398
siw events	2400
siw.cpubindingfailed events	2400
siw.mpaverrelaxed events	2400
siw.mpaversionmismatch events	2401
siw.pollingenablefailed events	2401
sk events	2402
sk.hog events	2402
sk.kt events	2402
sk.panic events	2402
sk.ring events	2403
sla events	2404
sla.shelf events	2404
sm events	2405
sm.break events	2405
sm.c2c events	2405
sm.cascade events	2406
sm.disable events	2407
sm.dst events	2408
sm.fgmirror events	2408
sm.keepcnt events	2409
sm.log events	2409

sm.mediator events	2410
sm.no events	2416
sm.slc events	2416
sm.syncmirror events	2417
sm.vfiler events	2417
sm.xfer events	2418
smbc events	2422
smbc.aufo events	2422
smbc.avail events	2424
smbc.cg events	2426
smbc.pfo events	2427
smbc.policy events	2429
smc events	2430
smc.cg events	2430
smc.dst events	2430
smc.frest events	2431
smc.restore events	2431
smc.snap events	2432
smc.snapmir events	2432
smc.snapmirror events	2438
smc.tot events	2439
smc.vlt events	2439
sms events	2441
sms.cg events	2441
sms.clone events	2441
sms.common events	2442
sms.failed events	2442
sms.fanout events	2443
sms.giveback events	2443
sms.operation events	2444
sms.qr events	2444
sms.resync events	2446
sms.snap events	2446
sms.src events	2447
sms.status events	2447
sms.transition events	2449
smtape events	2450
smtape.bkp events	2450
smtape.rst events	2452
snapdiff events	2455
snapdiff.rebl events	2455
snapdiff.spinnp events	2456
snapdiff.zapi events	2457
snaplock events	2458
snaplock.disk events	2458

snaplock.event events	2459
snaplock.invalid events	2460
snaplock.legal events	2460
snaplock.mcc events	2462
snaplock.nse events	2462
snaplock.sfo events	2463
snaplock.ss events	2463
snaplock.state events	2468
snaplock.sys events	2468
snaplock.trans events	2469
snaplock.user events	2470
snaplock.vcc events	2470
snaplock.vol events	2471
snaplock.volume events	2471
snapmirror events	2473
snapmirror.block events	2473
snapmirror.conf events	2474
snapmirror.copy events	2475
snapmirror.corrupted events	2477
snapmirror.dst events	2478
snapmirror.flextotraderr events	2495
snapmirror.log events	2495
snapmirror.lrs events	2496
snapmirror.metadata events	2497
snapmirror.migrate events	2498
snapmirror.multipath events	2502
snapmirror.reconstruct events	2502
snapmirror.registry events	2502
snapmirror.replay events	2503
snapmirror.resync events	2503
snapmirror.retry events	2504
snapmirror.sockreaderr events	2504
snapmirror.sockstreaderr events	2505
snapmirror.sockstrwriteerr events	2505
snapmirror.src events	2506
snapmirror.status events	2521
snapmirror.tradtoflexerr events	2522
snapmirror.unsupportedchksumtype events	2523
snapmirror_copy events	2524
snapmirror_copy.updatetimeincreased events	2524
snaprestore events	2525
snaprestore.fail events	2525
snaprestore.log events	2525
snaptags events	2526
snaptags.file events	2526

snaptags.snap events	2526
snapvault events	2527
snapvault.dst events	2527
snapvault.qtree events	2529
snapvault.reg events	2531
snapvault.src events	2531
snapvault.sys events	2532
snapvault.tgt events	2533
snmp events	2536
snmp.authentication events	2536
snmp.coldstart events	2536
snmp.fips events	2536
snmp.link events	2537
snmp.raid events	2538
snmp.server events	2539
snmp.snmpv3 events	2539
snmp.traphost events	2540
snmp.warmstart events	2540
socket events	2541
socket.limit events	2541
socket.sndbuflimit events	2541
sp events	2542
sp.auth events	2542
sp.authlog events	2542
sp.autoconfig events	2542
sp.autoupdate events	2543
sp.device events	2543
sp.dhcp events	2544
sp.dup events	2544
sp.eth events	2545
sp.firmware events	2546
sp.heartbeat events	2548
sp.ipmi events	2549
sp.log events	2549
sp.network events	2550
sp.notconfigured events	2552
sp.orftp events	2552
sp.postupd events	2553
sp.reboot events	2553
sp.servprocd events	2554
sp.snmp events	2554
sp.ssh events	2555
sp.upd events	2555
sp.userlist events	2556
sp.wdog events	2557

space events	2558
space.ratchet events	2558
spider events	2559
spider.message events	2559
splitter events	2560
splitter.owm events	2560
splog events	2561
splog.failed events	2561
splog.running events	2561
splog.sp events	2562
splog.unsupported events	2563
splog.warnings events	2563
spm events	2565
spm.avc events	2565
spm.bgpd events	2565
spm.coresegd events	2566
spm.cphmd events	2566
spm.crs events	2566
spm.cshmd events	2567
spm.enclmgmtd events	2567
spm.envmgr events	2568
spm.fpolicy events	2568
spm.hashd events	2569
spm.httprd events	2570
spm.ktlsd events	2570
spm.mes events	2570
spm.mgwd events	2571
spm.mntsvc events	2572
spm.named events	2572
spm.nchmd events	2573
spm.ndmpd events	2573
spm.nphmd events	2574
spm.ntpd events	2574
spm.perfstatd events	2575
spm.pipd events	2575
spm.schmd events	2576
spm.secd events	2576
spm.servprocd events	2577
spm.shmd events	2577
spm.timestate events	2577
spm.ucoreman events	2578
spm.upgrademgr events	2579
spm.vifmgr events	2579
spm.vldb events	2580
spm.vserverdr events	2580



spm.yplib events	2581
spmgmt events	2582
spmgmt.driver events	2582
sshd events	2584
sshd.auth events	2584
sshd.logingracetime events	2584
stackmon events	2585
stackmon.stack events	2585
storage events	2586
storage.cmd events	2586
storage.interface events	2586
storencr events	2588
storencr.disk events	2588
storvsc events	2589
storvsc.srb events	2589
suspend events	2590
suspend.zapi events	2590
sv events	2591
sv.vfiler events	2591
sv.vol events	2591
sync events	2593
sync.snapmirror events	2593
sysconfig events	2594
sysconfig.sysconfigtab events	2594
sysmgr events	2595
sysmgr.ecv events	2595
sysmgr.gb events	2595
sysmgr.upload events	2596
tape events	2600
tape.cmd events	2600
tape.config events	2602
tape.device events	2605
tape.diagmsg events	2607
tapemc events	2609
tapemc.alias events	2609
tapemc.device events	2610
tapemc.iopending events	2614
tapemc.passthru events	2615
tapemc.sas events	2615
tar events	2617
tar.csum events	2617
tar.extract events	2619
tcp events	2621
tcp.cc events	2621
tcp.change events	2621

tcpip events	2622
tcpip.conn events	2622
tcpip.duplicate events	2622
telnet events	2623
telnet.enabled events	2623
telnet.firewall events	2623
telnet.setsockopt events	2623
telnet.socket events	2624
telnet.threadcreate events	2625
tls events	2626
tls.insecure events	2626
tls.unused events	2627
token events	2628
token.copy events	2628
token.inactive events	2628
token.max events	2628
token.node events	2629
token.pit events	2630
token.vol events	2630
tsse events	2633
tsse.scan events	2633
ucm events	2635
ucm.disabled events	2635
ucm.mode events	2635
ucm.type events	2636
ucore events	2638
ucore.error events	2638
ucore.limit events	2638
ucore.limits events	2638
ucore.not events	2639
ucore.panicstring events	2639
ucore.ulimit0 events	2640
unowned events	2641
unowned.disk events	2641
upgrademgr events	2642
upgrademgr.cmd events	2642
upgrademgr.fw events	2642
upgrademgr.node events	2644
upgrademgr.package events	2644
upgrademgr.two events	2645
upgrademgr.update events	2645
usb events	2648
usb.device events	2648
usb.port events	2649
usbmon events	2651

usbmon.boot events	2651
usbmon.unable events	2651
user events	2653
user.group events	2653
useradmin events	2654
useradmin.added events	2654
useradmin.lockedout events	2654
useradmin.name events	2655
useradmin.registry events	2655
useradmin.rootaccess events	2656
vdisk events	2657
vdisk.clonelogfull events	2657
vdom events	2658
vdom.aggr events	2658
vdom.unmapped events	2658
vf events	2659
vf.cleanup events	2659
vfiler events	2660
vfiler.unable events	2660
vfilermigrate events	2661
vfilermigrate.diskownershipmismatch events	2661
vifmgr events	2662
vifmgr.bcastdomain events	2662
vifmgr.bd events	2662
vifmgr.bgp events	2663
vifmgr.cdb events	2666
vifmgr.clntdev events	2667
vifmgr.clus events	2668
vifmgr.cluscheck events	2668
vifmgr.clusterbd events	2675
vifmgr.clusterlfoconf events	2675
vifmgr.dbase events	2676
vifmgr.displaced events	2676
vifmgr.extroute events	2677
vifmgr.gcp events	2681
vifmgr.hm events	2682
vifmgr.hv events	2682
vifmgr.ifgrp events	2684
vifmgr.ipspace events	2684
vifmgr.lif events	2685
vifmgr.lifadded events	2690
vifmgr.lifbeinghosted events	2690
vifmgr.lifbeingremoved events	2690
vifmgr.lifcapacity events	2691
vifmgr.lifdeleted events	2691

vifmgr.lifdown events	2692
vifmgr.lifidoutofsync events	2692
vifmgr.lifmove events	2693
vifmgr.lifmoved events	2693
vifmgr.lifondegport events	2695
vifmgr.lifs events	2696
vifmgr.lifsuccessfullymoved events	2698
vifmgr.migratelifs events	2698
vifmgr.port events	2699
vifmgr.portdeg events	2707
vifmgr.portdown events	2709
vifmgr.porthhealthy events	2709
vifmgr.portowner events	2710
vifmgr.portremoved events	2710
vifmgr.portspeed events	2712
vifmgr.portup events	2712
vifmgr.rdma events	2713
vifmgr.reach events	2713
vifmgr.rpc events	2717
vifmgr.secondaryoffline events	2717
vifmgr.skippingl2pingtest events	2718
vifmgr.started events	2718
vifmgr.startup events	2719
vifmgr.staticroute events	2720
vifmgr.subnet events	2720
vifmgr.svcpolicy events	2722
vifmgr.unabletohostmovedlif events	2723
vifmgr.viplifwithoutbgp events	2723
vifmgr.vs events	2724
virtual events	2725
virtual.limits events	2725
vldb events	2726
vldb.adtcons events	2726
vldb.aggrbladeid events	2726
vldb.aggregatename events	2727
vldb.aggregateuuid events	2727
vldb.aggrupdate events	2728
vldb.aggrvote events	2729
vldb.mccaggregate events	2729
vldb.update events	2730
vldb.vldbagggr events	2730
vldb.vsrootvol events	2731
vm events	2732
vm.dev events	2732
vm.host events	2732

vnm events	2734
vnm.dataip events	2734
vnvram events	2735
vnvram.dma events	2735
vnvram.flush events	2735
vnvram.flushed events	2735
vol events	2737
vol.clone events	2737
vol.duplicate events	2737
vol.inconsistent events	2738
vol.language events	2738
vol.log events	2739
vol.max events	2741
vol.mount events	2741
vol.phy events	2741
vol.phys events	2742
vol.registryget events	2743
vol.reserve events	2744
volaggr events	2745
volaggr.offline events	2745
volcmd events	2746
volcmd.snapmirror events	2746
volcopy events	2747
volcopy.dst events	2747
volcopy.flextraderr events	2749
volcopy.src events	2749
volcopy.tradtoflexerr events	2750
volcreate events	2751
volcreate.fsid events	2751
volcreate.name events	2751
volmigrate events	2752
volmigrate.error events	2752
volmigrate.migrating events	2752
volmigrate.result events	2752
volume events	2754
volume.check events	2754
volume.transition events	2756
vpeer events	2757
vpeer.name events	2757
vpeer.request events	2757
vs events	2759
vs.migrate events	2759
vs.multiple events	2768
vsa events	2770
vsa.azure events	2770

vsa.cloud events	2771
vsa.device events	2774
vsa.diskobjectstore events	2774
vsa.ephemeral events	2775
vsa.instancemetadata events	2775
vsa.mlx events	2776
vsa.nodesernum events	2776
vsa.savecore events	2777
vsa.scheduledevent events	2778
vsa.vsphere events	2779
vscan events	2781
vscan.disabled events	2781
vscan.enabled events	2781
vscan.newversion events	2781
vscan.odjob events	2782
vscan.odreport events	2783
vscan.pool events	2784
vscan.privsharecreate events	2785
vscan.rdbupdregister events	2786
vsdr events	2787
vsdr.adtcreate events	2787
vsdr.app events	2787
vsdr.clone events	2788
vsdr.clonetmp events	2789
vsdr.cluster events	2789
vsdr.fpvollallpolicy events	2790
vsdr.fpvollallpolicyonfpaggr events	2790
vsdr.fpvoldropqospolicy events	2791
vsdr.fpvolonnonfpaggr events	2791
vsdr.invalid events	2792
vsdr.ip events	2792
vsdr.missing events	2793
vsdr.net events	2794
vsdr.nonfpvolonfpaggr events	2798
vsdr.periodic events	2799
vsdr.qos events	2800
vsdr.rootvol events	2801
vsdr.ruleaddedtosmpolicy events	2801
vsdr.rulemodifiedofsmpolicy events	2802
vsdr.schedule events	2802
vsdr.vol events	2803
vsdr.volmov events	2812
vsdr.volmove events	2812
vsdr.volmovsrcfpnonfpaggr events	2813
vsdr.volume events	2813

vserver events	2816
vserver.config events	2816
vserver.ipspace events	2816
vserver.mcc events	2818
vserver.name events	2819
vserver.oper events	2820
vserver.peering events	2820
vserver.rbac events	2822
vserver.rename events	2823
vserver.rootvolume events	2824
vserver.start events	2825
vserver.stop events	2826
vserver.storage events	2826
vserver.vsd events	2829
wackiron events	2830
wackiron.near events	2830
wackiron.panic events	2831
wackiron.past events	2831
waf events	2833
waf.32bit events	2833
waf.8 events	2833
waf.access events	2834
waf.afs events	2834
waf.aggr events	2835
waf.aggregate events	2844
waf.aggrtrans events	2845
waf.air events	2850
waf.analytics events	2850
waf.ap events	2856
waf.arl events	2857
waf.asserts events	2857
waf.atime events	2858
waf.attr events	2859
waf.bad events	2860
waf.bagofbits events	2861
waf.blkreclaim events	2862
waf.br events	2864
waf.buf events	2865
waf.c2c events	2866
waf.ca events	2869
waf.catalog events	2881
waf.cbin events	2882
waf.cde events	2884
waf.cdemap events	2884
waf.check events	2884

wafl.cleanup events	2885
wafl.clone events	2886
wafl.compaction events	2890
wafl.compress events	2890
wafl.convert events	2893
wafl.cop events	2894
wafl.cp events	2895
wafl.csm events	2896
wafl.data events	2896
wafl.delete events	2896
wafl.descend events	2898
wafl.dir events	2899
wafl.dirent events	2906
wafl.disable events	2906
wafl.dosname events	2907
wafl.enc events	2908
wafl.encrypted events	2908
wafl.encryption events	2909
wafl.err events	2910
wafl.extent events	2910
wafl.failed events	2911
wafl.fc events	2912
wafl.fgmember events	2912
wafl.filecg events	2913
wafl.fs events	2913
wafl.gd events	2915
wafl.giveback events	2916
wafl.gs events	2918
wafl.hya events	2920
wafl.i2p events	2924
wafl.ialloc events	2925
wafl.inaccessible events	2926
wafl.incomplete events	2926
wafl.incons events	2927
wafl.inconsistent events	2930
wafl.ino events	2936
wafl.inode events	2940
wafl.inofile events	2942
wafl.iron events	2942
wafl.latest events	2969
wafl.lmgr events	2970
wafl.lookup events	2970
wafl.make events	2971
wafl.maxdirsize events	2972
wafl.mcc events	2972



waf1.mem events	2979
waf1.memory events	2980
waf1.metadata events	2981
waf1.metafile events	2981
waf1.mirror events	2982
waf1.mount events	2983
waf1.mp events	2985
waf1.needs events	2985
waf1.nfsv4 events	2985
waf1.no events	2986
waf1.ntfsunixsecurityops events	2986
waf1.nvlog events	2987
waf1.nvsave events	2988
waf1.objstore events	2993
waf1.offline events	2993
waf1.ondemand events	2996
waf1.online events	3002
waf1.onoff events	3006
waf1.playlist events	3007
waf1.pnso events	3008
waf1.pvbn events	3008
waf1.qtree events	3010
waf1.quota events	3014
waf1.raid events	3023
waf1.rclm events	3030
waf1.readdir events	3030
waf1.reallocate events	3031
waf1.relocate events	3036
waf1.relocating events	3037
waf1.repl events	3038
waf1.replay events	3040
waf1.root events	3043
waf1.s2c events	3045
waf1.scan events	3046
waf1.seclabel events	3068
waf1.sfsr events	3068
waf1.shutdown events	3072
waf1.sidl events	3072
waf1.sis events	3074
waf1.slice events	3076
waf1.sm events	3077
waf1.snap events	3077
waf1.snaplock events	3082
waf1.snaprestore events	3083
waf1.snapshot events	3085

waf.snapshot_xlation_duplicate_entry events	3087
waf.snapshot_xlation_setupfailed events	3087
waf.spacemgmt events	3088
waf.spacetax events	3088
waf.sparse events	3089
waf.spcmgmt events	3089
waf.spcres events	3090
waf.spinhi events	3090
waf.split events	3091
waf.stream events	3091
waf.svo events	3092
waf.takeover events	3093
waf.timer events	3096
waf.trans events	3096
waf.transition events	3097
waf.unable events	3097
waf.upgrade events	3098
waf.vol events	3098
waf.volinfo events	3125
waf.volmove events	3125
waf.volref events	3126
waf.voltrans events	3126
waf.volume events	3132
waf.volume_snap events	3137
waf.vv events	3138
waf.vvclock events	3143
waf.vvol events	3144
waf.watch events	3153
waf.write events	3153
waf.zombie events	3155
wafiron events	3156
wafiron.summary events	3156
xf cvi events	3157
xf cvi.attr events	3157
xf cvi.cm events	3157
xf cvi.cross events	3158
xf cvi.disabled events	3158
xf cvi.fcrenabled events	3158
xf cvi.initfailed events	3159
xf cvi.invalid events	3159
xf cvi.link events	3160
xf cvi.qlgc events	3161
xf cvi.qlgc_regdumpfailed events	3178
xf cvi.unsupported events	3179
xf cvi.update events	3179

xinetd events	3180
xinetd.hit events	3180
xyz events	3182
xyz.trans events	3182
zapi events	3183
zapi.killed events	3183
zapi.retired events	3183
zapi.sf events	3183
zapi.snapshot events	3184
zrto events	3185
zrto.lif events	3185
zsm events	3186
zsm.send events	3186
zsm.socket events	3186
ztl events	3188
ztl.smap events	3188
Legal notices	3189
Copyright	3189
Trademarks	3189
Patents	3189
Privacy policy	3189
Open source	3189

# ONTAP 9.14.1 EMS reference

# accesscache events

## accesscache.cfg events

### accessCache.cfg.updateFailed

#### Severity

ERROR

#### Description

This message occurs when exports access cache configuration is modified but the new configuration could not be updated in the Nblade. The access cache is therefore operating with the older configuration.

#### Corrective Action

Retry the exports access cache configuration modification. If it continuously fails, restart management gateway. If the problem persists, contact NetApp technical support.

#### Syslog Message

Access cache configuration change for Vserver "%s" on node "%s" could not be updated. Failure reason is "%s".

#### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

**node** (STRING): Name of the Node on which the problem is occurring.

**reason** (STRING): The reason for the failure.

## accesscache.nearlimits events

### AccessCache.NearLimits

#### Severity

ALERT

#### Description

This message occurs when the access cache module is near its limits for entries or export rules. Reaching these limits can prevent new clients from being able to mount and perform I/O on the storage system, and can also cause clients to be granted or denied access based on stale cached information.

#### Corrective Action

Ensure that the number of clients accessing the storage system continues to be below the limits for access cache entries and export rules across those entries. If the set of clients accessing the storage system is constantly changing, consider using the "vserver export-policy access-cache config modify" command to reduce the harvest timeout parameter so that cache entries for clients that are no longer accessing the storage system can be evicted sooner.

#### Syslog Message

The exports access cache module is near its limits for entries (current count : %u, warning threshold : %u, limit : %u) or export rules across entries (current count : %u, warning threshold : %u, limit : %u).

## Parameters

**numEntriesCurrent** (INT): Current number of entries in the cache.  
**numEntriesWarningThreshold** (INT): Warning threshold for number of entries.  
**numEntriesLimit** (INT): Limit on the number of entries.  
**numRulesCurrent** (INT): Current number of export rules in the cache.  
**numRulesWarningThreshold** (INT): Warning threshold for number of export rules.  
**numRulesLimit** (INT): Limit on the number of export rules.

# accesscache.reachedlimits events

## AccessCache.ReachedLimits

### Severity

EMERGENCY

### Description

This message occurs when the access cache module reaches its limit for entries or export rules. This condition can prevent new clients from being able to mount and perform I/O on the storage system. This condition can also cause clients to be granted or denied access based on stale information.

### Corrective Action

Ensure that the number of clients accessing the storage system is below the limits for access cache entries and export rules across those entries. If the set of clients accessing the storage system is constantly changing, consider using the "vserver export-policy access-cache config modify" command to reduce the harvest timeout parameter so that cache entries for clients that are no longer accessing the storage system can be evicted sooner.

### Syslog Message

The exports access cache module on this node reached its limits while trying to add %d new entries (current count : %u, limit : %u) and/or %d export rules (current count : %u, limit : %u).

## Parameters

**numEntriesDelta** (INT): Number of new entries to be added.  
**numEntriesCurrent** (INT): Current number of entries in the cache.  
**numEntriesLimit** (INT): Limit on the number of entries.  
**numRulesDelta** (INT): Number of new export rules to be added.  
**numRulesCurrent** (INT): Current number of export rules in the cache.  
**numRulesLimit** (INT): Limit on the number of export rules.

# acp events

## acp.auto events

### acp.auto.upgrade

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the ACP Processor has an older firmware image compared to the file revision present in '/etc/acp/acpp\_fw'. The ACP Administrator is performing an automatic firmware upgrade.

#### Corrective Action

(None).

#### Syslog Message

Performing automatic firmware upgrade (image: %s) on %s (%s).

#### Parameters

**file** (STRING): Command issued.

**inband\_id** (STRING): Inband ID of the module.

**ip** (STRING): IP address of the module.

## acp.bind events

### acp.bind.error

#### Severity

ERROR

#### Description

This message occurs when the system encounters a socket binding error while setting up the Alternate Control Path (ACP) Management Ethernet port. ACP functionality is inactive after encountering this.

#### Corrective Action

Verify that the Ethernet cable is properly attached to the ACP port and the green LED is on. Reenable the ACP port using the "options acp.enabled" command.

#### Syslog Message

An error was encountered while configuring the ACP Management Ethernet Port: %s.

#### Parameters

**ErrorString** (STRING): Error string.

## acp.commerr events

## acp.commErr

### Severity

ALERT

### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the SAS shelf I/O module (IOM) that is configured for in-band ACP is experiencing communication errors.

### Corrective Action

Examine the event log for I/O errors that might have a common source such as faulty cabling, and correct them. If the errors persist, the IOM might need to be replaced. Contact NetApp technical support.

### Syslog Message

IOM %s is reporting multiple in-band ACP communication errors. Shelf serial number: %s.

### Parameters

**inband\_id** (STRING): Inband ID of the module.

**shelf\_sno** (STRING): Storage shelf serial number of the shelf where the I/O module is seated.

## acp.commerrpersist events

### acp.commErrPersist

### Severity

ALERT

### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that a SAS shelf I/O module (IOM) configured for in-band ACP is experiencing persistent communication errors.

### Corrective Action

Examine the event log for I/O errors that might have a common source with this problem, such as faulty cabling. If other devices are not experiencing failures, then contact NetApp technical support to replace the IOM.

### Syslog Message

IOM %s has failed after reporting multiple in-band ACP errors. Shelf serial number: %s.

### Parameters

**inband\_id** (STRING): Inband ID of the module.

**shelf\_sno** (STRING): Storage shelf serial number of the shelf where the IOM is seated.

## acp.dev events

### acp.dev.notify

### Severity

ERROR



## Description

This message occurs when the Alternate Control Path (ACP) Administrator encounters an error event in a development-only kernel.

## Corrective Action

Notify ACP development. No user action is needed.

## Syslog Message

ACP error: %s

## Parameters

**debug\_string** (STRING): Error statement that occurred in a development kernel.

# acp.downrev events

## acp.downrev.acpp

### Severity

ERROR

### Description

This message occurs when the Alternate Control Path Administrator cannot successfully download up-to-date firmware to an Alternate Control Path Processor (ACPP) or when the most up-to-date firmware image is not present in the `/etc/acpp_fw/` directory.

### Corrective Action

Verify that the latest ACPP firmware is available in the firmware directory. If the firmware update continually fails to download up-to-date firmware, then replace the ACPP I/O module.

### Syslog Message

ACPP module %s (%s) has downrev firmware.

### Parameters

**inband\_id** (STRING): In-band ID of the ACPP module.

**ip** (STRING): IP address of the ACPP module.

# acp.exp events

## acp.exp.power.cycle

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Control Path (ACP) Processor receives a command to power-cycle the SAS expander.

### Corrective Action

(None).

## Syslog Message

SAS expander power-cycle request sent to %s (%s), (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Status code returned by the 'expander power cycle' command.

## acp.exp.power.cycle.failed

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Controller Path Administrator fails in its attempt to power-cycle the SAS expander.

### Corrective Action

(None).

## Syslog Message

SAS expander power-cycle request sent to %s (%s) failed with status (%s), (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Reason for the 'expander power-cycle' failure.

**shelf\_sno** (STRING): Shelf serial number of the shelf where this I/O module is seated.

## acp.exp.power.cycle.success

### Severity

INFORMATIONAL

### Description

This message occurs when an 'expander power cycle' command sent to the Alternate Controller Path (ACP) Processor succeeds.

### Corrective Action

(None).

## Syslog Message

SAS expander power cycle command sent to %s (%s) was successful, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## **acp.exp.reset**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the Alternate Control Path (ACP) Processor receives a command to reset a SAS expander.

### **Corrective Action**

(None).

### **Syslog Message**

SAS expander reset issued to %s (%s), (disk shelf serial number: %s).

### **Parameters**

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Status code returned by the 'expander reset' command.

## **acp.exp.reset.failed**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the Alternate Controller Path Administrator fails in its attempt to reset the SAS expander.

### **Corrective Action**

(None).

### **Syslog Message**

SAS expander reset request sent to %s (%s) failed with status (%s), (disk shelf serial number: %s).

### **Parameters**

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Reason for the 'expander reset' failure.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## **acp.exp.reset.success**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an 'expander reset' command sent to the Alternate Controller Path (ACP) Processor succeeds.

## Corrective Action

(None).

## Syslog Message

SAS expander reset command sent to %s (%s) was successful, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.exp.vpd.failed

### Severity

INFORMATIONAL

### Description

This message occurs when an attempt to get SAS expander VPD (Vital Product Data) data from the Alternate Control Path Processor fails.

### Corrective Action

(None).

### Syslog Message

An 'expander vpd' request sent to %s (%s) failed with status (%s), (disk shelf serial number: %s).

### Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Reason for the failure to get the SAS expander vpd data.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.fw events

### acp.fw.download

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Alternate Control Path (ACP) Processor receives a command to initiate a firmware download.

#### Corrective Action

(None).

#### Syslog Message

Firmware image %s will be downloaded to %s (%s), (disk shelf serial number: %s).

## Parameters

**command** (STRING): Firmware image name.  
**inband\_id** (STRING): In-band ID of the module.  
**ip** (STRING): IP address of the module.  
**status** (STRING): Status code returned by the 'initiation' command.

## acp.fw.download.failed

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Controller Path Administrator fails in its attempt to download firmware.

### Corrective Action

(None).

### Syslog Message

Firmware download on %s (%s) failed with status (%s), (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.  
**ip** (STRING): IP address of the module.  
**status** (STRING): Type of failure that occurred during the firmware download.  
**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.fw.download.success

### Severity

INFORMATIONAL

### Description

This message occurs when a 'firmware download' command sent to the Alternate Controller Path Processor succeeds.

### Corrective Action

(None).

### Syslog Message

Firmware download on %s (%s) was successful, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): In-band ID of the module.  
**ip** (STRING): IP address of the module.  
**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.init events

## acp.init.configIP

### Severity

ERROR

### Description

This message occurs when ONTAP® software cannot configure the Ethernet port designated for the Alternate Control Path (ACP) with 30 seconds, and marks the ACP as inactive.

### Corrective Action

Verify the connection to the disk shelf ACP port, and make sure that it is not connected to a public network. Try plugging it into another Ethernet port with no devices.

### Syslog Message

Could not configure the ACP Ethernet port within 30 seconds.

### Parameters

(None).

## acp.init.netmask

### Severity

INFORMATIONAL

### Description

This message occurs during initialization state, when the Alternate Control Path (ACP) Administrator netmask was configured incorrectly. The system sets the netmask to a default value.

### Corrective Action

(None).

### Syslog Message

Setting ACP Administrator netmask value to %s (reason: %s).

### Parameters

**value** (STRING): Default netmask.

**reason** (STRING): Reason to choose the default.

## acp.init.ssl

### Deprecated

Deprecated as of ONTAP® version 9.2. The ACP process was never really responsible for initializing SSL, and has not tried to do so for a very long time.

### Severity

ERROR

### Description

This message occurs when the Alternate Control Path (ACP) process cannot initialize the underlying Secure Socket Library (SSL). The ACP Administrator will be inactive.

**Corrective Action**

(None).

**Syslog Message**

Could not initialize SSL.

**Parameters**

(None).

## acp.invalid events

### acp.invalid.vpd.info

**Severity**

ALERT

**Description**

This message occurs when the Alternate Control Path (ACP) system detects that an Alternate Control Path Processor (ACPP) has invalid VPD (Vital Product Data) info.

**Corrective Action**

Replace the Alternate Control Path Processor (ACPP).

**Syslog Message**

ACPP module %s (%s) has invalid VPD.

**Parameters**

**inband\_id** (STRING): In-band ID of the ACPP module.

**ip** (STRING): IP address of the ACPP module.

## acp.ioxm events

### acp.ioxm.port.down

**Severity**

INFORMATIONAL

**Description**

This message occurs when the link status of one or more ports from the IOXM switch changes from "up" to "down".

**Corrective Action**

Verify that all cables in the IOXM switch ports are seated properly.

**Syslog Message**

IOXM switch port %d is down.

**Parameters**

**port\_num** (INT): Switch port number of the IOXM switch.

## acp.ioxm.port.up

### Severity

INFORMATIONAL

### Description

This message occurs when the link status of one or more ports from the IOXM switch changes from "down" to "up".

### Corrective Action

(None).

### Syslog Message

IOXM switch port %d is up.

### Parameters

**port\_num** (INT): Switch port number of the IOXM switch.

## acp.locked events

### acp.locked.wrench.link.down

### Severity

INFORMATIONAL

### Description

This message occurs when the link status of a locked wrench Ethernet port (on the back of the storage controller) changes from "up" to "down".

### Corrective Action

Verify that the cable in the port is seated properly.

### Syslog Message

The on-board locked wrench port is down.

### Parameters

(None).

### acp.locked.wrench.port.down

### Severity

ERROR

### Description

This message occurs when the link status of a locked wrench Ethernet port (on the back of the storage controller) changes from "up" to "down".

### Corrective Action

Verify that the cable in the port is seated properly.



## Syslog Message

The on-board locked wrench port is down; Alternate Control Path (ACP) management is inactive.

## Parameters

(None).

## acp.locked.wrench.port.up

## Severity

INFORMATIONAL

## Description

This message occurs when the link status of a locked wrench Ethernet port (on the back of the storage controller) changes from "down" to "up".

## Corrective Action

(None).

## Syslog Message

The on-board locked wrench port is up.

## Parameters

(None).

## acp.need events

## acp.need.firmware.update

## Severity

ERROR

## Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the ACP processor is running an older firmware revision that supports only the Secure Socket Library (SSL) channel. The "/etc/acpp\_fw" directory requires the latest firmware file for the update to succeed. This module is marked as inactive.

## Corrective Action

Copy the latest firmware file for this module into "/etc/acpp\_fw", and then use the "system node run -node <nodename> -command storage download acp" command to download the firmware.

## Syslog Message

ACP module %s (%s) needs a firmware upgrade, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): Inband ID of the module.

**ip** (STRING): IP address of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.need.module.upgrade

### Severity

ERROR

### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the I/O module (IOM) does not support InBand ACP.

### Corrective Action

Upgrade the IOM to IOM6 or later.

### Syslog Message

IOM %s does not support InBand ACP. Upgrade the IOM to IOM6 or later for supportability (disk shelf serial number: %s).

### Parameters

**inband\_id** (STRING): Inband ID of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this IOM is seated.

## acp.no events

### acp.no.ip

### Severity

ERROR

### Description

This message occurs when the node cannot find a free IP address for the Alternate Control Path (ACP) Management port. ACP functionality is disabled.

### Corrective Action

Check the connection to the disk shelf ACP port and verify that it is not connected to a public network. Try plugging it into another Ethernet port with no devices, to see whether it is changed.

### Syslog Message

Could not assign a free IP address to the ACP Management Ethernet port, disabling ACP functionality.

### Parameters

(None).

## acp.post events

### acp.post.data.failed

### Severity

INFORMATIONAL

### Description

This message occurs when a 'post data' request that was sent to the Alternate Controller Path Processor

fails.

### Corrective Action

(None).

### Syslog Message

A 'post data' request sent to %s (%s) failed with status (%s), (disk shelf serial number: %s).

### Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Reason for the 'post data' failure.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

## acp.shelf events

### acp.shelf.power.cycle.failed

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Alternate Control Path (ACP) Processor fails to complete a power cycle of the disk shelf.

#### Corrective Action

(None).

#### Syslog Message

Disk shelf power cycle request issued to %s failed because %s.

#### Parameters

**shelf\_id** (STRING): In-band ID or Shelf Serial Number of the module.

**status** (STRING): Reason for failure.

### acp.shelf.power.cycle.success

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Alternate Control Path (ACP) Processor successfully completes a power cycle of the disk shelf.

#### Corrective Action

(None).

#### Syslog Message

Disk shelf power cycle request issued to %s is successful.

## Parameters

**shelf\_id** (STRING): In-band ID or Shelf Serial Number of the module.

## acp.shelf.power.off

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Control Path (ACP) Processor receives a command to power off the disk shelf.

### Corrective Action

(None).

### Syslog Message

Disk shelf power off request issued to %s

### Parameters

**shelf\_id** (STRING): In-band ID of the module or Shelf Serial Number.

## acp.shelf.power.off.failed

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Controller Path Administrator fails in its attempt to power-off the SAS disk shelf.

### Corrective Action

(None).

### Syslog Message

Disk shelf power off request sent to %s failed with status %s

### Parameters

**shelf\_id** (STRING): In-band ID of the module.

**status** (STRING): Reason for the power-off failure.

## acp.shelf.power.off.success

### Severity

INFORMATIONAL

### Description

This message occurs when a 'shelf power off' command sent to the Alternate Controller Path (ACP) Processor succeeds.

**Corrective Action**

(None).

**Syslog Message**

Disk shelf power off command sent to %s was successful

**Parameters**

**shelf\_id** (STRING): In-band ID of the module or Shelf Serial Number.

**acp.shelf.power.on****Severity**

INFORMATIONAL

**Description**

This message occurs when the Alternate Control Path (ACP) Processor receives a command to power on a disk shelf.

**Corrective Action**

(None).

**Syslog Message**

Disk shelf power up request issued to %s

**Parameters**

**shelf\_id** (STRING): In-band ID of the module or Shelf Serial Number.

**acp.shelf.power.on.failed****Severity**

INFORMATIONAL

**Description**

This message occurs when the Alternate Controller Path Administrator fails in its attempt to power on the SAS disk shelf.

**Corrective Action**

(None).

**Syslog Message**

Disk shelf power on request sent to %s failed with status %s

**Parameters**

**shelf\_id** (STRING): In-band ID or Shelf Serial Number of the module.

**status** (STRING): Reason for the 'power-on' failure.

**acp.shelf.power.on.success****Severity**

INFORMATIONAL

## Description

This message occurs when a 'shelf power on' command sent to the Alternate Controller Path (ACP) Processor succeeds.

## Corrective Action

(None).

## Syslog Message

Disk shelf power on command sent to %s was successful

## Parameters

**shelf\_id** (STRING): In-band ID of the module or Shelf Serial Number.

# acp.sock events

## acp.sock.error

### Severity

ERROR

### Description

This message occurs when the system encounters a socket error while setting up the Alternate Control Path (ACP) Management Ethernet port. ACP functionality is subsequently inactive.

### Corrective Action

Verify that the Ethernet cable is properly attached to the ACP port and that the green LED is on. Reenable the ACP using the "options acp.enabled on" command.

### Syslog Message

An error was encountered while configuring the ACP Management Ethernet Port: %s.

### Parameters

**ErrorString** (STRING): Error string.

# acp.update events

## acp.update.failed

### Severity

ERROR

### Description

This message occurs when the Alternate Control Path (ACP) Administrator receives notification from the ACP Processor that it cannot upgrade the firmware because it is in some faulty state or the firmware file is not valid. The ACP processor will continue to run with the older firmware version.

### Corrective Action

Try issuing the firmware download again using the "system node run -node <nodename> -command storage download acp" command.

## Syslog Message

ACP module %s (%s) failed to upgrade firmware, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): Inband ID of the module.

**ip** (STRING): IP address of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

# acp.updateconfiguration events

## acp.updateConfiguration

### Severity

ALERT

### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the SAS shelf I/O module (IOM) supports only the in-band ACP configuration.

### Corrective Action

> Use the "storage shelf acp configure -channel in-band -is-enabled true" command to enable in-band ACP.

## Syslog Message

IOM %s in shelf S/N: %s does not support out-of-band ACP.

## Parameters

**inband\_id** (STRING): Inband ID of the module.

**shelf\_sno** (STRING): Storage shelf serial number of the shelf where the IOM is seated.

# acp.upgrade events

## acp.upgrade.successful

### Severity

INFORMATIONAL

### Description

This message occurs when the Alternate Control Path (ACP) Administrator determines that the ACP Processor has successfully upgraded the firmware.

### Corrective Action

(None).

## Syslog Message

ACP module %s (%s) successfully upgraded firmware, (disk shelf serial number: %s).

## Parameters

**inband\_id** (STRING): Inband ID of the module.

**ip** (STRING): IP address of the module.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.

# acp.voltage events

## acp.voltage.status.failed

### Severity

INFORMATIONAL

### Description

This message occurs when a 'voltage status' request that was sent to the Alternate Controller Path Processor fails.

### Corrective Action

(None).

### Syslog Message

A voltage status request sent to %s (%s) failed with status (%s), (disk shelf serial number: %s).

### Parameters

**inband\_id** (STRING): In-band ID of the module.

**ip** (STRING): IP address of the module.

**status** (STRING): Reason for the 'voltage status' failure.

**shelf\_sno** (STRING): Disk shelf serial number of the shelf where this I/O module is seated.



# adminapi events

## adminapi.input events

### adminapi.input.badXml

**Severity**

ERROR

**Description**

This message occurs when the Data ONTAP server receives a bad XML request that does not contain the API name to be executed.

**Corrective Action**

Correct the input XML with the API name to be executed.

**Syslog Message**

Input XML does not contain the API name to be executed.

**Parameters**

(None).

## adminapi.xml events

### adminapi.xml.encoding

**Severity**

ERROR

**Description**

This event is generated in case of a mismatched XML encoding received by ONTAP server.

**Corrective Action**

Correct the mismatched encoding and retry the API.

**Syslog Message**

ONTAP server got unexpected encoding (expected %s, got %s

**Parameters**

**expEnco** (STRING): The expected xml encoding.

**recvEnco** (STRING): The actual received encoding.

### adminapi.xml.namespace

**Severity**

ERROR

**Description**

This event is generated in case of a mismatched xml namespace received by ONTAP server.

**Corrective Action**

Correct the XML namespace and retry the API.

**Syslog Message**

ONTAP server namespace (expected %s, got %s

**Parameters**

**expNamesp** (STRING): The expected xml namespace.

**uri** (STRING): The received uri

# adt events

## adt.aggrcb events

### adt.aggrcb.unlinkfail

#### Severity

ALERT

#### Description

This message occurs when the auditing subsystem cannot remove a stale aggregate control block entry from the Audit aggregate hash table. This could happen if the hash table rearrangement fails because of insufficient memory. In this condition, the auditing subsystem might not function or data access might be denied.

#### Corrective Action

Reboot the node. For high-availability (HA) configurations, perform a takeover and giveback from the partner to reduce the impact on system availability

#### Syslog Message

Auditing subsystem internal error: Aggregate Control Block: %s unlink failed for SVM with UUID: %s because of insufficient memory.

#### Parameters

**AggregateUuid** (STRING): UUID of the aggregate with a stale file handle.

**VserverUuid** (STRING): UUID of the Storage Virtual Machine(SVM, formerly Vserver).

## adt.aucb events

### adt.aucb.overwrite

#### Severity

INFORMATIONAL

#### Description

This message occurs in the auditing subsystem, when the subsystem overwrite an existing audit control block during audit control block creation. The overwriting of the audit control block is because of a failure in audit control block deletion.

#### Corrective Action

(None).

#### Syslog Message

Overwriting of audit control block occurred.

#### Parameters

(None).

# adt.cmdq events

## adt.cmdq.init.fail

### Severity

ALERT

### Description

This message occurs when an internal audit command queue fails to initialize.

### Corrective Action

Available system resources are required for this feature to function. The audit subsystem might not function, or might function with reduced performance. The best way to accomplish that is through a system reboot. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. After the reboot is complete, issue a "giveback" command to return services to this partner.

### Syslog Message

Audit subsystem internal error: Audit command queue failed to initialize, system reboot might be required.

### Parameters

(None).

# adt.cmdthrd events

## adt.cmdthrd.create.fail

### Severity

ALERT

### Description

This message occurs when audit initialization fails because the necessary execution threads cannot be created.

### Corrective Action

Available system resources are required for this feature to function. The best way to accomplish that is through a system reboot. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. After the reboot is complete, issue a "giveback" command to return services to this partner.

### Syslog Message

Audit subsystem internal error: Could not start required audit command threads, system reboot might be required.

### Parameters

(None).

# adt.dest events

## adt.dest.access.fail

### Severity

ERROR

### Description

This message occurs when the audit consolidation job fails to write audit logs to the destination directory because the destination directory is not available. This event can lead to denial of service on Security Access Control List (SACL) enabled objects.

### Corrective Action

Use the "volume show" command to check whether the destination volume exists, is online, and has sufficient free space. Use the "volume modify" command to bring the volume online, or increase volume size. Use the "vserver audit modify" command to modify the destination volume.

### Syslog Message

Audit consolidation job failed to access destination directory "%s" of Vserver "%s"; consolidation job type is "%s".

### Parameters

**destination** (STRING): Destination directory of the Vserver into which the consolidation job failed to write audit logs.

**vserver** (STRING): Vserver name whose consolidation job failed to write audit logs to the destination directory.

**descriptor** (STRING): Type of consolidation job - CIFS/S3.

## adt.dest.directory.full

### Severity

EMERGENCY

### Description

This message occurs when the audit consolidation process fails to write audit logs to the destination directory because the destination directory is full. This event can lead to denial of service on Security Access Control List (SACL) enabled objects.

### Corrective Action

Use the "volume show" command to check whether the destination volume exists, is online, and has sufficient free space. Use the "volume modify" command to increase volume size. Use the "vserver audit modify" command to modify the destination volume path.

### Syslog Message

Audit destination directory "%s" of Vserver "%s" is full; consolidation job type is "%s".

### Parameters

**destination** (STRING): Destination directory of the Vserver that is full.

**vserver** (STRING): Vserver name whose audit consolidation process failed to write audit logs to the destination directory.

**descriptor** (STRING): Type of consolidation job - CIFS/S3.

## adt.dest.directory.unavail

### Severity

EMERGENCY

### Description

This message occurs when the audit consolidation job fails to write audit logs to the destination directory because the destination directory has become unavailable or has run out of space. This event can lead to denial of service for Security Access Control List (SACL) enabled objects.

### Corrective Action

Use the "volume show" command to check whether the destination volume exists, is online, and has sufficient free space. Use the "volume modify" command to bring the volume online, or increase volume size. Use the "vserver audit modify" command to modify the destination volume.

### Syslog Message

Audit destination directory "%s" of Vserver "%s" is unavailable or out of space; consolidation job type is "%s".

### Parameters

**destination** (STRING): Destination directory of the Vserver to which the consolidation job failed to write audit logs.

**vserver** (STRING): Name of the Vserver name whose consolidation job failed to write audit logs to the destination directory.

**descriptor** (STRING): Type of consolidation job - CIFS/S3.

## adt.max events

### adt.max.record.size.exceeded

### Severity

ERROR

### Description

This message occurs when a management command or file operation tries to generate an audit record, that is greater than the max\_audit\_record\_size value. The File Services Auditing subsystem might fail the operation or truncate the record.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

File Services Auditing subsystem failed the operation or truncated an audit record because it was greater than max\_audit\_record\_size value. Vserver UUID=%s, event\_id=%u, size=%u

### Parameters

**VserverUuid** (STRING): UUID of then Vserver.

**eventId** (INT): Event ID for internal use.

**size** (INT): Record size.

# adt.nfs events

## adt.nfs.root.access.no.map

### Severity

ERROR

### Description

This message occurs in the auditing subsystem, when the NFS root user is accessing the file without a valid CIFS credential. Access is being allowed because the "ignore-nt-acl-for-root" NFS option is set, and hence the BUILTIN\Administrator credential is used for auditing.

### Corrective Action

If NFS root user access to the file without NT ACL check should be prohibited, disable the "ignore-nt-acl-for-root" option by using the "nfs modify -ignore-nt-acl-for-root" command. If a more appropriate CIFS credential should be used for auditing, configure the SVM name-mapping by using the "vserver name-mapping" command.

### Syslog Message

NFS root user is accessing the file without a valid CIFS credential. Access is being allowed because the "ignore-nt-acl-for-root" NFS option is set, and hence the BUILTIN\Administrator credential is used for auditing.

### Parameters

**VserverUuid** (STRING): UUID of the Storage Virtual Machine (SVM, formerly Vserver).

# adt.service events

## adt.service.block

### Severity

EMERGENCY

### Description

This message occurs when the volume that hosts the destination path where the final audit logs are stored is not accessible or does not have sufficient space and the audit consolidation service is not able to write to the volume. This prevents the audit service from working properly. File access and file audit logging is blocked.

### Corrective Action

Use the "vserver audit show" command to get the name of the volume that hosts the destination path. Use the "volume show" command to check whether the volume exists, is online, and has sufficient free space. Use the "volume modify" command to increase volume size. Use the "vserver audit modify" command to modify the destination path.

### Syslog Message

Audit destination directory "%s" of Vserver "%s" is either not accessible or full for the consolidation job type "%s". File-Operations which need to be audited will be denied. File Audit Logging will also fail.

## Parameters

**destinationPath** (STRING): Destination path.

**vserverName** (STRING): Name of the Storage Virtual Machine (SVM, formerly Vserver).

**descriptor** (STRING): Type of consolidation job - CIFS/S3.

## adt.service.ro.filesystem

### Severity

EMERGENCY

### Description

This message occurs when the volume that hosts the destination path where the final audit logs are stored is of type Data protection (DP) or Load sharing (LS). These volume types are read-only for the purposes of audit logging and logging will fail.

### Corrective Action

Use the "vserver audit modify" command to modify the destination volume.

### Syslog Message

Auditing and file-ops failing on SVM %s, the volume containing %s for consolidation job of type %s is read only filesystem.

## Parameters

**vserverName** (STRING): Name of the Storage Virtual Machine (SVM, formerly Vserver).

**destinationPath** (STRING): Destination path.

**descriptor** (STRING): Type of consolidation job - CIFS/S3.

## adt.service.unblock

### Severity

INFORMATIONAL

### Description

This message occurs when the volume that hosts the destination path where the final audit logs are stored was resized or cleaned up and the audit consolidation service is able to write to the volume thereby unblocking the audit service. File access and file audit logging is unblocked.

### Corrective Action

(None).

### Syslog Message

Audit destination directory "%s" of Vserver "%s" for consolidation job of type "%s" is now accessible and File-Operations which need to be audited will be unblocked. File Audit Logging will also be unblocked.

## Parameters

**destinationPath** (STRING): Destination path.

**vserverName** (STRING): Name of the Storage Virtual Machine (SVM, formerly Vserver).

**descriptor** (STRING): Type of consolidation job - CIFS/S3.



# adt.stgvol events

## adt.stgvol.info.notavailable

### Severity

ERROR

### Description

This message occurs when retrieving staging volume information fails. This might be due to an internal system problem or the staging volume creation failed.

### Corrective Action

Contact Cserver admin to check if staging volume has been created.

### Syslog Message

Audit subsystem internal error: Either staging volume %s is not created or unable to retrieve staging volume information.

### Parameters

**volName** (STRING): The Name of the staging volume.

## adt.stgvol.missing

### Severity

ALERT

### Description

This message occurs when retrieving staging volume information fails. Either the staging volume was not created or information cannot be retrieved due to some internal error. As a result, auditing is not performed for operations on the corresponding volume.

### Corrective Action

Use the command "vserver audit repair" in diagnostic mode to correct the problem.

### Syslog Message

Audit subsystem internal error: Staging volume "%s" is either not created or not available.

### Parameters

**volName** (STRING): Name of the volume that is missing.

## adt.stgvol.nospace

### Severity

EMERGENCY

### Description

This message occurs when a staging volume does not have enough space and the audit service tries to create a new staging file or directory for writing audit logs.

### **Corrective Action**

Check the volume for older staging files that are no longer required and remove those files. Then retry the operation.

### **Syslog Message**

Audit subsystem internal error: Staging volume %s is full.

### **Parameters**

**volName** (STRING): Name of the volume that is full.

## **adt.stgvol.offline**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a staging volume is offline and the audit service tries to write an audit log into the staging volume.

### **Corrective Action**

Contact the Cserver administrator to set the staging volume status to online.

### **Syslog Message**

Audit subsystem internal error: Staging volume %s is offline.

### **Parameters**

**volName** (STRING): The Name of the volume which is offline.

## **adt.wrkrq events**

### **adt.wrkrq.init.fail**

### **Severity**

ALERT

### **Description**

This message occurs when an internal audit worker queue fails to initialize.

### **Corrective Action**

Available system resources are required for this feature to function. The Audit subsystem might not function, or might function with reduced performance. The best way to accomplish that is through a system reboot. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. After the reboot is complete, issue a "giveback" command to return services to this partner.

### **Syslog Message**

Audit subsystem internal error: Audit worker queue failed to initialize; system reboot might be required.

### **Parameters**

(None).

# adt.wrkrthrd events

## adt.wrkrthrd.create.fail

### Severity

ALERT

### Description

This message occurs when audit initialization fails because the necessary execution threads cannot be created.

### Corrective Action

Available system resources are required for this feature to function. The best way to accomplish that is through a system reboot. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. After the reboot is complete, issue a "giveback" command to return services to this partner.

### Syslog Message

Audit subsystem internal error: Could not start required audit worker threads; system reboot might be required.

### Parameters

(None).

# aggr events

## aggr.full events

### aggr.full

#### Severity

ALERT

#### Description

This message occurs when an aggregate is full.

#### Corrective Action

Ensure that all containers in the pod are online and healthy. Check if the REST server is up and responding.

#### Syslog Message

Aggregate "%s" is full.

#### Parameters

**aggr\_name** (STRING): Name of the aggregate.

## aggr.grow events

### aggr.grow.completed

#### Severity

NOTICE

#### Description

This message occurs when request to grow an aggregate is completed successfully. The request to grow the aggregate could be DMAP-triggered or SDE-triggered.

#### Corrective Action

(None).

#### Syslog Message

Request to grow aggregate "%s" is complete. Old size: %lu, new size: %lu, old state: %d, new state: %d.

#### Parameters

**aggr\_name** (STRING): Name of the aggregate.

**old\_size** (LONGINT): Size of the aggregate before it is grown.

**new\_size** (LONGINT): Size of the aggregate after it is grown.

**previous\_state** (INT): State of the aggregate before it is grown.

**current\_state** (INT): State of the aggregate after it is grown.

## aggr.inconsistent events

## aggr.inconsistent.modify

### Severity

INFORMATIONAL

### Description

This message occurs when the 'aggr inconsistent' command is used to modify a state of an aggregate.

### Corrective Action

(None).

### Syslog Message

Aggregate %s: 'aggr inconsistent' command used to change inconsistent state to '%s'.

### Parameters

**aggr** (STRING): Name of the aggregate.

**state** (STRING): State of the action: "set" or "clear".

## aggr.max events

### aggr.max.size.increased

### Severity

INFORMATIONAL

### Description

This message occurs as a reminder that the maximum aggregate size limit is increased from the default value defined in the Platform Specific Module (PSM) table.

### Corrective Action

(None).

### Syslog Message

Maximum aggregate size limit is increased from %ld TB to %ld TB.

### Parameters

**old\_limit** (LONGINT): Original aggregate size limit as defined in the PSM table.

**new\_limit** (LONGINT): New maximum aggregate size limit.

## aggr.provisioning events

### aggr.provisioning.unassigned.detected

### Severity

NOTICE

### Description

This message is issued when the aggregate provisioning recommendation generated for the cluster is found to be incomplete due to the presence of unassigned disks. Any spare disks that are not owned by a node in the cluster cannot contribute to the usable capacity of new aggregates.

## Corrective Action

Use the following steps to resolve the issue: - Run 'storage disk show -container-type unassigned' to list all the unassigned disks in the cluster. To assign all unassigned disks, run 'storage disk assign -auto true -node [node-name]'. - If there are still unassigned disks in the cluster, refer to the Storage Management Guide for manual disk assignment best practices. - For further information or assistance, contact NetApp technical support.

## Syslog Message

The generated aggr provisioning recommendation might be incomplete because of the following reason: %s

## Parameters

**details** (STRING): Details of the number of unassigned disks visible to a given node

**source** (STRING): The source that generated this warning.

# aggr.relocation events

## aggr.relocation.failed

### Severity

ERROR

### Description

This message occurs when relocation of an aggregate fails.

### Corrective Action

Refer to other generated error messages to identify the exact cause of failure and the corrective action.

### Syslog Message

Relocation of aggregate %s failed by node %s because %s.

### Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.

**node** (STRING): Name of the node that failed the relocation.

**msg** (STRING): Error that prevented relocation.

# aggr.resize events

## aggr.resize.request.sent

### Severity

NOTICE

### Description

This message occurs when a request to resize an aggregate is sent to the REST server as part of the storage thin provisioning feature in Quark.

### Corrective Action

(None).

## Syslog Message

Triggering resizing of aggregate for %s. Current size: %lu, new size: %lu.

## Parameters

**aggr\_name** (STRING): Name of the aggregate.

**current\_size** (LONGINT): Current size of the aggregate.

**new\_size** (LONGINT): New size of the aggregate after being resized.

## aggr.resize.response.failure

### Severity

ALERT

### Description

This message occurs when the REST server returns an error for a request to resize an aggregate.

### Corrective Action

Ensure that all containers in the pod are online and healthy. Check if the REST server is up and responding.

## Syslog Message

Server returned error: %d for the resizing request. Aggregate: %s, current size: %lu, new size: %lu".

## Parameters

**error** (INT): Error returned from REST server.

**aggr\_name** (STRING): Name of the aggregate.

**current\_size** (LONGINT): Current size of the aggregate.

**new\_size** (LONGINT): New size of the aggregate after being resized.

## aggr.resize.response.success

### Severity

NOTICE

### Description

This message occurs when the REST server returns "success" as a response to the request to resize an aggregate.

### Corrective Action

(None).

## Syslog Message

Server returned "success" as a response to a resizing request. Aggregate: %s, current size: %lu, new size: %lu.

## Parameters

**aggr\_name** (STRING): Name of the aggregate.

**current\_size** (LONGINT): Current size of the aggregate.

**new\_size** (LONGINT): New size of the aggregate after being resized.

# aggr.wafliron events

## aggr.wafliron.attemptStart

### Severity

INFORMATIONAL

### Description

This message occurs when wafliron is started. A later message indicates the success or failure of wafliron.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**volname** (STRING): Volume name.

**arguments** (STRING): Arguments supplied to wafliron on the command line.



# aggrcopy events

## aggrcopy.dst events

### aggrcopy.dst.autoRestrictMsg

**Severity**

NOTICE

**Description**

This event is issued when an aggregate copy start is issued on an aggregate which is offline.

**Corrective Action**

Aggr copy start is issued on a restricted aggregate use "aggr restrict" to restrict the destination aggregate.

**Syslog Message**

Aggr copy automatically restricting destination aggregate %s (currently offline).

**Parameters**

**dst** (STRING): The aggr copy destination aggregate

### aggrcopy.dst.noMemory

**Severity**

ERROR

**Description**

This event is issued when there is insufficient memory on the aggr copy destination for the transfer to continue.

**Corrective Action**

(None).

**Syslog Message**

System is low on memory, aborting aggr copy destination transfer.

**Parameters**

(None).

### aggrcopy.dst.revert

**Severity**

NOTICE

**Description**

This event is issued when an aggr copy start operation has to revert the checksum state of the destination volume aggregate because the source and destination aggregates have differing checksum states.

## Corrective Action

No corrective action required.

## Syslog Message

An aggr copy transfer to destination aggregate %s has found that the source and destination aggregates have differing checksum states. Data ONTAP is aborting the transfer and reverting the checksum state for the destination aggregate. After Data ONTAP displays a message showing the revert is complete, you will need to retry the operation. You can use the "aggr status -c" command to see the progress of the reversion.

## Parameters

**dst** (STRING): The aggr copy destination volume

## aggrcopy.dst.revertErr

### Severity

NOTICE

### Description

This event is issued when a aggr copy start operation is not able to revert the checksum state of the destination aggregate.

### Corrective Action

(None).

## Syslog Message

Data ONTAP cannot revert the checksum state of destination aggregate %s.

## Parameters

**dst** (STRING): The aggr copy destination

## aggrcopy.dst.reverting

### Severity

NOTICE

### Description

This event is issued when an aggr copy start operation is issued on an aggregate which is reverting its checksum state.

### Corrective Action

(None).

## Syslog Message

Aborting aggr copy transfer to aggregate %s because it is reverting its checksum state.

## Parameters

**dst** (STRING): The aggr copy destination aggregate

# aggrcopy.restart events

## aggrcopy.restart.dst.chkFail

### Severity

NOTICE

### Description

This message occurs when 'aggr copy' checks for restart capability on the destination storage system fail. The 'aggr copy' transfer operation cannot be restarted. It starts as usual from the beginning (0 %) without any further user action.

### Corrective Action

No action necessary

### Syslog Message

Destination side checks fail; 'aggr copy' transfer starts from 0%%.

### Parameters

(None).

# aggrcopy.src events

## aggrcopy.src.snapshotErr

### Severity

ERROR

### Description

This event is issued when we are not able to obtain all the snapshot information, from the vol copy source volume, required for a vol copy transfer.

### Corrective Action

(None).

### Syslog Message

Aggr copy not able to obtain snapshot information from source aggregate %s

### Parameters

**srcaggr** (STRING): The aggr copy source volume

# air events

## air.certify events

### air.certify.cancel

#### Severity

ALERT

#### Description

This message occurs when a metadata inconsistency is discovered inside a file system feature that is read-only and therefore cannot be immediately corrected. One or more processes in ONTAP® software might have been canceled when they encountered the inconsistency.

#### Corrective Action

No immediate corrective action is necessary, because the ONTAP process that encountered the inconsistency has been automatically canceled. If the condition persists, contact NetApp technical support for assistance.

#### Syslog Message

AIR certification of %s failed: %s, %s ; message %s canceled for subtype %d

#### Parameters

**fh** (STRING): File handle of the object with an inconsistency.  
**condition** (STRING): Specific metadata validation test that failed.  
**source** (STRING): Location of the failing metadata validation test.  
**msg** (STRING): The ONTAP process that encountered the inconsistency.  
**subtype** (INT): Subtype of the inode encountering the inconsistency.

## air.cleared events

### air.cleared.qtm.entry

#### Severity

NOTICE

#### Description

This message occurs when the Automated Incremental Repair (AIR) subsystem clears an entry in the qtree metafile because of inconsistent data within that entry.

#### Corrective Action

(None).

#### Syslog Message

Qtree entry certification failed at %s. In the qtree metafile %d, AIR cleared an entry that corresponds to the TID %d in volume %s.

#### Parameters

**source** (STRING): Location where qtree entry certification failed.  
**inode** (INT): Inode number of qtree metafile.

**tid** (INT): TID.

**volume** (STRING): Name of the volume.

## air.corrected events

### air.corrected.qtree.tid

#### Severity

NOTICE

#### Description

This message occurs when the Automated Incremental Repair (AIR) subsystem repairs an inode having a corrupted qtree ID.

#### Corrective Action

(None).

#### Syslog Message

Qtree entry certification failed at %s. AIR repaired qtree inode %d: old TID %d, new TID %d, fgindex %d.

#### Parameters

**source** (STRING): Location where qtree entry certification failed.

**inode** (INT): Inode number.

**old\_tid** (INT): Old TID.

**new\_tid** (INT): New TID.

**fgindex** (INT): Fgindex of the qtree root.

## air.disable events

### air.disable.async.delete.cli

#### Severity

ALERT

#### Description

This message occurs when any inconsistency is found with the async delete trashbin directory of a volume.

#### Corrective Action

The "asynchronous directory delete from the client" feature has been disabled on this volume. To re-enable the feature, use the (privilege: advanced) "volume file async-delete client enable" command

#### Syslog Message

The "async directory delete from the client" feature was disabled in volume %s, after an inconsistency was detected with the trashbin directory.

#### Parameters

**volume** (STRING): Name of the volume.

**source** (STRING): Location of the failing metadata validation.

## air.disable.preserve.unlink

### Severity

NOTICE

### Description

This message occurs when any inconsistency is found with the trash directory that is used for the preserve unlink option of a volume. Then, the preserve unlink option is disabled for that volume. The trash directory that is used for the preserve unlink option is used to preserve deleted files for NFSv41 when they have existing share locks.

### Corrective Action

None. The system reverts to the behavior it used before the preserve unlink option was enabled; that is, it will not preserve deleted files with share locks.

### Syslog Message

The preserve unlink option of volume "%s", was disabled after an inconsistency with the trash directory was detected.

### Parameters

**volume** (STRING): Name of the volume.  
**source** (STRING): Location of the failing metadata validation.

## air.enabled events

### air.enabled.fg.qtree

### Severity

NOTICE

### Description

This message occurs in a FlexGroup, when the qtree support is enabled on the volume, but disabled on one or more constituent volumes. ONTAP® software automatically enabled the qtree support on the relevant constituent volumes.

### Corrective Action

This message indicates that the qtree support is enabled on the volume, but disabled in one or more constituents within the FlexGroup. ONTAP® software automatically enabled the qtree support in the relevant constituent FlexGroup volumes.

### Syslog Message

Enabled qtree support in FlexGroup constituent volume %s after detecting inconsistency in %s.

### Parameters

**volume** (STRING): Name of the volume.  
**source** (STRING): Location of the failing metadata validation.

## air.evicted events

## air.evicted.qtm.entry

### Severity

NOTICE

### Description

This message occurs when the Automated Incremental Repair (AIR) subsystem evicts an entry in the qtree metafile because of inconsistent data within that entry.

### Corrective Action

(None).

### Syslog Message

Qtree entry certification failed at %s. In the qtree metafile %d, AIR evicted an entry that corresponds to the TID %d in volume %s.

### Parameters

**source** (STRING): Location where qtree entry certification failed.

**inode** (INT): Inode number of the qtree metafile.

**tid** (INT): TID.

**volume** (STRING): Name of the volume.

## air.fc events

### air.fc.origin.corrupt

#### Severity

ALERT

#### Description

This message occurs when inconsistencies are detected in the origin of this FlexCache® volume. The affected files remain inaccessible from this FlexCache volume until the corrective action is taken.

#### Corrective Action

If the origin of this FlexCache volume is a FlexVol® volume, then run wafiron on the origin FlexVol volume.

If the origin is a FlexGroup volume, then list the affected files from a NFS or a CIFS client on the origin volume. For more information or assistance, contact NetApp technical support.

#### Syslog Message

The fh %s in Origin of this FlexCache volume %s%s contains inconsistencies.

#### Parameters

**fh** (STRING): File handle of the affected inode.

**vol** (STRING): FlexCache volume name.

**volident** (STRING): Unique identifier for the volume if volume name alone is insufficient.

## air.qtm events

## air.qtm.rebuild.scan.state

### Severity

NOTICE

### Description

This message occurs when the state of a scanner used to rebuild the qtree metafile changes.

### Corrective Action

(None).

### Syslog Message

Scanner to rebuild the qtree metafile is %s.

### Parameters

**state** (STRING): State of the scanner.

## air.removed events

### air.removed.remote.lck.entry

### Severity

NOTICE

### Description

This message occurs when there is an inconsistency in the remote lock entry metafile. The inconsistent entry is removed upon detection.

### Corrective Action

(None).

### Syslog Message

Remote lock entry %s certification failed: expression %s, source %s, state (%s).

### Parameters

**entry** (STRING): Lock entry type and origin file ID pair describing the entry.

**rlem** (STRING): Describes the metafile to which this lock entry belongs to.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

**state** (STRING): Description of the remote lock entry metafile metadata state that was found to be inconsistent.

## air.repair events

### air.repair.prsist.lock.entry

### Severity

NOTICE



## Description

This message occurs when there is an inconsistency in the persistent lock entry metafile. The inconsistent entry is repaired upon detection if there is a corresponding in-core record.

## Corrective Action

(None).

## Syslog Message

Persistent lock entry %s certification failed: Condition: %s, source: %s, state: %s.

## Parameters

**entry** (STRING): Lock entry type and origin file ID pair describing the entry.

**condition** (STRING): Initial validation test that fails.

**source** (STRING): Location of the failing metadata validation test.

**state** (STRING): Description of the remote lock entry metafile metadata state that is found to be inconsistent.

# air.repaired.events

## air.repaired.bucket.header

### Severity

NOTICE

### Description

This message occurs when ONTAP® discovers an inconsistency in a metadata bucket header record. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

TOC metadata (%s) has an inconsistent record for the bucket upload table (%s). %s.

### Parameters

**toc\_fh** (STRING): File handle of the TOC metadata file.

**upload\_table\_fh** (STRING): File handle of the bucket upload table file.

**state** (STRING): Description of the metadata that was found to be inconsistent.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.cbmap

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the type/subtype of CloudBlockMap metafile. The inconsistency is automatically repaired immediately upon detection.

**Corrective Action**

(None).

**Syslog Message**

CloudBlockMap metafile %s certification failed: %s, %s (%s)

**Parameters**

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the CloudBlockMap metafile that was found to be inconsistent.

**air.repaired.cbmap.entry****Severity**

NOTICE

**Description**

This message occurs when context mismatch is discovered in any entry of CloudBlockMap metafile. The inconsistency is automatically repaired immediately upon detection.

**Corrective Action**

(None).

**Syslog Message**

%s certification failed: %s, %s (%s)

**Parameters**

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the metafile that was found to be inconsistent.

**pvbn** (LONGINT): Inconsistency found in cbmap entry of this PVBN.

**air.repaired.cbmap.header****Severity**

NOTICE

**Description**

This message occurs when an inconsistency is discovered in the header of CloudBlockMap metafile. The inconsistency is automatically repaired immediately upon detection.

**Corrective Action**

(None).

**Syslog Message**

%s certification failed: %s, %s (%s)

**Parameters**

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.  
**source** (STRING): The location of the failing metadata validation test.  
**state** (STRING): A description of the metafile that was found to be inconsistent.

## air.repaired.chapter.entry

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in a metadata record. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Chapter metadata (%s) has inconsistent record for object (%s). Source %s (%s).

### Parameters

**chapter\_fh** (STRING): File handle of the metadata file.  
**object\_fh** (STRING): File handle of the object.  
**source** (STRING): Location of the failing metadata validation test.  
**condition** (STRING): Initial metadata validation test that failed.  
**state** (STRING): Description of the metadata that was found to be inconsistent.

## air.repaired.dir.hole

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the hole list of a directory. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

Directory hole list %s certification failed: %s, %s.

### Parameters

**fh** (STRING): File handle of the affected directory.  
**condition** (STRING): Initial metadata validation test that failed.  
**source** (STRING): Location of the failing metadata validation test.

## air.repaired.dir.inode

### Severity

NOTICE

## Description

This message occurs when an inconsistency is discovered in a directory. The inconsistency is repaired automatically as soon as it is detected.

## Corrective Action

(None).

## Syslog Message

Directory %s certification failed: %s, %s.

## Parameters

**fh** (STRING): File handle of the affected directory.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.enc.vvol.info

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the encryptable aggregate metafile inodes encrypted flexvol information which stores vvol\_bt看id and vvol\_fid. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

Inode %s certification failed: %s, %s.

### Parameters

**fh** (STRING): File handle of the affected inode.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.fabriclink

### Severity

NOTICE

### Description

This message occurs when a metadata inconsistency is discovered and repaired within an object storage replication data structure.

### Corrective Action

(None).

### Syslog Message

AIR repaired %s, condition %s at %s, object "%s"

## Parameters

**fh** (STRING): File handle of the replication data structure that was found to be inconsistent.  
**expr** (STRING): The nature of the inconsistency that was discovered and repaired.  
**source** (STRING): Location where the metadata inconsistency was detected.  
**object** (STRING): Optionally represents the name of a particular object whose replication state was affected.

## air.repaired.fc.dir

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in any directory in the volume. The inconsistency is repaired automatically as soon as it is discovered.

### Corrective Action

(None).

### Syslog Message

Evicted inconsistent directory with fh %s due to corruption in volume %s.

## Parameters

**fh** (STRING): File handle of the affected inode.  
**volume** (STRING): Name of the volume.

## air.repaired.fc.qtree.id

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the qtree ID of an inode. The inconsistent inode is evicted automatically as soon as it is discovered.

### Corrective Action

(None).

### Syslog Message

Evicted inconsistent inode with fh %s due to a corrupted qtree ID: qtree ID %d, volume %s.

## Parameters

**fh** (STRING): File handle of the affected inode.  
**tid** (INT): Qtree ID.  
**volume** (STRING): Name of the volume.

## air.repaired.hardlink.i2p

### Severity

NOTICE

## Description

This message occurs when the AIR subsystem repairs an inconsistent inode and previously unreachable hard links are returned to availability.

## Corrective Action

(None).

## Syslog Message

AIR repaired inconsistent hardlink for inode %d: %d %d %s

## Parameters

**inode** (INT): Inode number.  
**parent\_inode** (INT): Parent inode number.  
**link\_count** (INT): Inode link count.  
**volume** (STRING): Name of the volume.

## air.repaired.label.database

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered within a WAFL label database. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Label database %s certification failed: %s, %s (%s)

### Parameters

**fh** (STRING): File handle of the affected metafile.  
**condition** (STRING): The initial metadata validation test that failed.  
**source** (STRING): The location of the failing metadata validation test.  
**state** (STRING): A description of the metadata that was found to be inconsistent.

## air.repaired.lsm.bfc.lkp

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistency in the LSM BFC Lookup file. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

LSM BFC Lookup file %s was corrected at block %d. Source %s.

## Parameters

**fh** (STRING): File handle of the affected file.  
**block** (LONGINT): Block number of the inconsistent block.  
**source** (STRING): Location of the failing metadata validation test.

## air.repaired.lsm.bloom.file

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistency in the LSM Bloomfilter file. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

LSM Bloomfilter file %s was corrected at block %d. Source %s.

## Parameters

**fh** (STRING): File handle of the affected file.  
**block** (LONGINT): Block number of the inconsistent block.  
**source** (STRING): Location of the failing metadata validation test.

## air.repaired.lsm.info.file

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the LSM Info file. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

LSM Info file %s was corrected at block %d. Source %s.

## Parameters

**fh** (STRING): File handle of the affected file.  
**block** (INT): Block number of the inconsistent block.  
**source** (STRING): Location of the failing metadata validation test.

## air.repaired.lsm.keyspc.file

### Severity

NOTICE

## Description

This message occurs when the system detects an inconsistency in the LSM Keyspace Information file. The inconsistency is repaired automatically as soon as it is detected.

## Corrective Action

(None).

## Syslog Message

Inconsistent LSM Keyspace Information metafile on aggregate "%s" was repaired. Source %s.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.lsm.lookup.file

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistency in the LSM Lookup file. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

LSM Lookup file %s was corrected at block %d. Source %s.

### Parameters

**fh** (STRING): File handle of the affected file.

**block** (LONGINT): Block number of the inconsistent block.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.merkle

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the metadata (Merkle tree file). Repairs to the Merkle tree file have automatically been started.

### Corrective Action

(None).

### Syslog Message

Merkle file "%s" certification failed: %s, directory ID (%llu), branch ID (%llu), buffer level %d, fbn (%llu), Merkle tree level (%d)



## Parameters

**fh** (STRING): File handle of the Merkle file that needs repairs.  
**condition** (STRING): Initial validation test that failed.  
**dir\_id** (LONGINT): Slice directory ID of the Merkle file.  
**branch\_id** (LONGINT): Slice branch ID of the Merkle file.  
**level** (INT): Buffer level of the Merkle file block.  
**fbn** (LONGINT): File block number (FBN) of the Merkle file.  
**merkle\_level** (INT): Merkle checksum level of the Merkle file.  
**source** (STRING): The location of the failing validation test.

## air.repaired.mpu.catalog

### Severity

NOTICE

### Description

This message occurs when ONTAP® discovers an inconsistency in a multipart upload catalog record of a bucket upload table. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Bucket upload table (%s) has an inconsistent record for the multipart upload catalog (%s). %s.

## Parameters

**upload\_table\_fh** (STRING): File handle of the bucket upload table file.  
**mp\_catalog\_fh** (STRING): File handle of the multipart catalog file.  
**state** (STRING): Description of the metadata that was found to be inconsistent.  
**condition** (STRING): Initial metadata validation test that failed.  
**source** (STRING): Location of the failing metadata validation test.

## air.repaired.multipart.database

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the metadata of a multipart inode. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Multipart inode %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected file.  
**condition** (STRING): The database validation test that failed.  
**source** (STRING): The location of the failing database validation test.  
**state** (STRING): A description of the metadata that was found to be inconsistent.

## air.repaired.multipart.inode

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in a multipart inode. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Multipart inode %s certification failed: %s, %s

### Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): The initial multipart inode validation test that failed.

**source** (STRING): The location of the failing multipart inode validation test.

## air.repaired.multipart.record

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the metadata of the multipart inode. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Multipart inode %s certification failed: %s, %s (%s)

### Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): The initial record validation test that failed.

**source** (STRING): The location of the failing record validation test.

**state** (STRING): A description of the record that was found to be inconsistent.

## air.repaired.multipart.rectify

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the rectification metadata for a multipart file. The inconsistency is automatically repaired immediately upon detection.

## Corrective Action

(None).

## Syslog Message

Multipart inode %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): The initial record validation test that failed.

**source** (STRING): The location of the failing record validation test.

**state** (STRING): A description of the record that was found to be inconsistent.

## air.repaired.multipart.userheader

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the userheader section of a multipart file. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Multipart inode %s userheader certification failed: %s, %s

### Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): The initial userheader validation test that failed.

**source** (STRING): The location of the failing userheader validation test.

## air.repaired.name.info.flags

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the name information flags of a directory. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

Name information flags for directory entry %s corrected at block %d and index %d. Source %s.

### Parameters

**fh** (STRING): File handle of the affected directory.

**block** (INT): Block number that contains the affected entry

**index** (INT): Location within the block that holds the affected entry

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.pct.entry

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in a metadata record. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

PCT metadata has inconsistent record for bucket (%u) file (%u). Source %s (%s, %s).

### Parameters

**bucketid** (INT): Bucket identifier.

**fileid** (INT): File identifier.

**source** (STRING): Location of the failing metadata validation test.

**condition** (STRING): Initial metadata validation test that failed.

**state** (STRING): Description of the metadata that was found to be inconsistent.

## air.repaired.qtree.id

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the qtree ID of an inode. The inconsistency is repaired automatically as soon as it is discovered.

### Corrective Action

(None).

### Syslog Message

Repaired inconsistent inode with fh %s due to a corrupted qtree ID: Old qtree ID %d, new qtree ID %d, volume %s.

### Parameters

**fh** (STRING): File handle of the affected inode.

**old\_tid** (INT): Old qtree ID.

**new\_tid** (INT): New qtree ID.

**volume** (STRING): Name of the volume.

## air.repaired.qtree.metafile

### Severity

NOTICE

**Description**

This message occurs when an inconsistency is discovered in a qtree metafile. The inconsistency is repaired automatically as soon as it is detected.

**Corrective Action**

(None).

**Syslog Message**

Qtree metafile %s

**Parameters**

**details** (STRING): Description of the metadata failure that has been fixed.

**air.repaired.qtree.root****Severity**

NOTICE

**Description**

This message occurs when the system discovers an inconsistency between the qtree file root and a qtree metafile entry. The system repairs this inconsistency automatically, as soon as it detects it.

**Corrective Action**

(None).

**Syslog Message**

Qtree root %s certification failed: %s, %s corrections: %s.

**Parameters**

**fh** (STRING): File handle of the affected qtree root.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

**state** (STRING): Description of the metadata that has been fixed.

**air.repaired.ral.clone.exp****Severity**

NOTICE

**Description**

This message occurs when an inconsistency is discovered between the source, clone created via ral retrieve and the clone expansion structure in source's remote state.

**Corrective Action**

(None).

**Syslog Message**

Remote RAL Clone %s certification failed: %s, %s (%s)

**Parameters**

**fh** (STRING): File handle of the affected clone.

**condition** (STRING): The initial metadata validation test that failed.  
**source** (STRING): The location of the failing metadata validation test.  
**state** (STRING): A description of the source and clone files.

## **air.repaired.rclone.database**

### **Severity**

NOTICE

### **Description**

This message occurs when an inconsistency is discovered within the WAFL Remote Clone Database. The inconsistency is automatically repaired immediately upon detection.

### **Corrective Action**

(None).

### **Syslog Message**

Remote Clone Database %s certification failed: %s, %s (%s)

### **Parameters**

**fh** (STRING): File handle of the affected metafile.  
**condition** (STRING): The initial metadata validation test that failed.  
**source** (STRING): The location of the failing metadata validation test.  
**state** (STRING): A description of the metafile that was found to be inconsistent.

## **air.repaired.rclone.record**

### **Severity**

NOTICE

### **Description**

This message occurs when an inconsistency is discovered in any record of the WAFL Remote Clone Database. The inconsistency is automatically repaired immediately upon detection.

### **Corrective Action**

(None).

### **Syslog Message**

Remote Clone Record %s certification failed: %s, %s (%s)

### **Parameters**

**fh** (STRING): File handle of the affected record.  
**condition** (STRING): The initial record validation test that failed.  
**source** (STRING): The location of the failing record validation test.  
**state** (STRING): A description of the entry that was found to be inconsistent.

## **air.repaired.remote.index**

### **Severity**

NOTICE

## Description

This message occurs when an inconsistency is discovered within the WAFL RAL index. The inconsistency is automatically repaired immediately upon detection.

## Corrective Action

(None).

## Syslog Message

Remote index %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): The RAL index record that was discovered to be incorrect.

## air.repaired.remote.inode

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered between a file and its RAL state within WAFL. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Remote inode %s certification failed: %s, %s (%s)

### Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the RAL state metadata that was found to be inconsistent.

## air.repaired.remote.metafile

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered within a WAFL RAL metafile. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Remote metafile %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the RAL metafile that was found to be inconsistent.

## air.repaired.remote.tallies

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the tallies in Remote Entry metafile. The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Remote metafile %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the RAL metafile that was found to be inconsistent.

## air.repaired.rlem.tallies

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the tallies in Remote Lock Entry metafile(RLEM). The inconsistency is automatically repaired immediately upon detection.

### Corrective Action

(None).

### Syslog Message

Remote lock entry metafile %s certification failed: %s, %s (%s)

## Parameters

**fh** (STRING): File handle of the affected metafile.

**rlem** (STRING): Describes the metafile to which this lock entry belongs to.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**state** (STRING): A description of the RLEM that was found to be inconsistent.



## air.repaired.slice

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in the metadata (slice file). Repairs to the slice file from the secondary copy have automatically been started.

### Corrective Action

(None).

### Syslog Message

Slice file "%s" certification failed: %s, dir ID (%llu), branch ID (%llu), level %d, fbn (%llu)

### Parameters

**fh** (STRING): File handle of the slice file that needs repairs.  
**condition** (STRING): The initial validation test that failed.  
**dir\_id** (LONGINT): Slice directory ID of the slice file.  
**branch\_id** (LONGINT): Slice branch ID of the slice file.  
**level** (INT): Buffer level of the slice file block.  
**fbn** (LONGINT): File block number (FBN) of the slice file.  
**source** (STRING): The location of the failing validation test.

## air.repaired.snapinfo.block

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistency in the persistent lock entry metafile. The inconsistent entry has been removed.

### Corrective Action

(None).

### Syslog Message

Repaired inconsistent block (%s) found in Snapinfo metafile on volume "%s".

### Parameters

**state** (STRING): Description of the Snapinfo metafile block that was found to be inconsistent.  
**volume** (STRING): Name of the volume.

## air.repaired.ssm.iftp.inode

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered in ssm iftp metafile. The inconsistency is repaired automatically as soon as it is detected.

### Corrective Action

(None).

### Syslog Message

AIR vlist metafile %s certification failed: %s, %s.

### Parameters

**fh** (STRING): File handle of the affected metafile.

**condition** (STRING): Initial metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.toc.entry

### Severity

ALERT

### Description

This message occurs when an inconsistency is discovered in a metadata record. The inconsistency is automatically repaired immediately upon detection. But it can leave behind lost objects that are no longer pointed to by the namespace.

### Corrective Action

The command "object-store-server lost-object-recovery start" under diag privilege can be used to add lost objects back into namespace. The progress of this command can be monitored using "object-store-server lost-object-recovery show" command.

### Syslog Message

TOC metadata (%s) has inconsistent record for chapter (%s). %s (%s).

### Parameters

**toc\_fh** (STRING): File handle of the TOC metadata file.

**chapter\_fh** (STRING): File handle of the chapter metadata file.

**state** (STRING): Description of the metadata that was found to be inconsistent.

**condition** (STRING): Initial metadata validation test that failed.

## air.repaired.user.ind

### Severity

ERROR

### Description

This message occurs when an inconsistency is discovered in the user indirect block. The inconsistency is repaired automatically immediately upon detection.

### Corrective Action

(None).

### Syslog Message

User indirect %s certification failed: %s (%s)

## Parameters

**fh** (STRING): File handle of the affected file.

**condition** (STRING): Initial metadata validation test that failed.

**state** (STRING): Description of the user indirect block that was found to be inconsistent.

**source** (STRING): Location of the failing metadata validation test.

## air.repaired.vplus

### Severity

NOTICE

### Description

This message occurs when an indexed storage tree is rebuilt to repair a metadata inconsistency that was discovered.

### Corrective Action

(None).

### Syslog Message

AIR vplus rebuild of %s completed for subtype %d

## Parameters

**fh** (STRING): File handle of the storage tree.

**subtype** (INT): Subtype of the inode encountering the inconsistency.

## air.vplus events

### air.vplus.certify.issue

#### Severity

NOTICE

#### Description

This message occurs when an inconsistency is discovered inside an indexed storage tree. The inconsistency is automatically repaired immediately upon detection.

#### Corrective Action

(None).

#### Syslog Message

AIR vplus certification of %s failed: %s, %s for subtype %d

## Parameters

**fh** (STRING): File handle of the storage tree.

**condition** (STRING): Specific metadata validation test that failed.

**source** (STRING): Location of the failing metadata validation test.

**subtype** (INT): Subtype of the inode encountering the inconsistency.

# am events

## am.dest events

### am.dest.check.failed

#### Severity

NOTICE

#### Description

Relocation of aggregates fails at the destination as we no longer accept aggregates from versions 8.2.0 or lower.

#### Corrective Action

User needs to upgrade to 8.2.1 before attempting ARL to 8.3.0 or higher. `

#### Syslog Message

Relocation of aggregates failed from node %s to current node as ARL is no longer supported from releases 8.2.0 or lower to 8.3.0 or higher.

#### Parameters

**partner\_node** (STRING): Source of the aggregate migration

# api events

## api.engine events

### api.engine.killed

#### Severity

EMERGENCY

#### Description

This event is posted when the ONTAPI® engine is unexpectedly terminated. The ONTAPI APIs cannot be executed until the system is rebooted.

#### Corrective Action

Reboot the controller. If the problem persists, contact Contact NetApp technical support.

#### Syslog Message

Error in API Engine. Engine name: %s.

#### Parameters

**engineName** (STRING): Name of the execution engine (e.g., "m" or "d").

## api.fileio events

### api.fileio

#### Severity

NOTICE

#### Description

This message occurs after a file I/O error that typically causes a Data ONTAP® API to fail. While processing Data ONTAP API requests, the system might write intermediate results to files. For example, a long running command returning quantities of data might send a batch of data and save remaining data in a file for subsequent delivery. An I/O error occurred during one of these operations.

#### Corrective Action

If the problem is related to creating a file, use the system's "df" and "df -i" commands to determine whether the root volume is full or all inodes are in use. If so, it might be possible to remove unneeded files to release space or inodes needed for Data ONTAP API operations. If the problem is that an intermediate file is being deleted because of its age, no action is needed. Probably the client that requested the API did not complete the API processing.

#### Syslog Message

(None).

#### Parameters

**errorDescription** (STRING): Description of the I/O error that occurred.

**targetDescription** (STRING): Description of the file or directory in which the error occurred.

**errorCode** (INT): Numeric error code.

**errorDetail** (STRING): Description of the error code.

# api.input events

## api.input.invalidSchema

### Severity

ERROR

### Description

This message occurs when an ONTAPI(tm) API receives extra input parameters as part of the API request or when a required input parameter is missing in the API request. Most often this is due to a bug in the application that initiates the API request.

### Corrective Action

The error is due to strict schema validation. You can turn off the input schema validation by using the following CLI commands: `$priv set advanced $registry set state.api.schema_input_validate.enable off`

### Syslog Message

Error in API input schema validation. API name: %s. Detailed error: %s.

### Parameters

**apiName** (STRING): Name of the ONTAPI API in which something invalid was detected.

**detailInfo** (STRING): Description of the error in the API input, such as an extra input value or a missing input value.

# api.output events

## api.output.invalid

### Severity

NOTICE

### Description

This message occurs when An invalid character was detected while checking the output for a Data ONTAP® API. All Data ONTAP API output must be valid UTF-8 characters. From time to time invalid output is detected. Most often this is due to hardware errors, for example, a disk-list-info API might query a disk drive for its serial number but receive invalid data from a malfunctioning drive.

### Corrective Action

Examine the information in detailInfo to see if it pinpoints defective hardware. If an API name is provided, that API can be executed to determine which field in the API output is being given invalid data.

### Syslog Message

(None).

### Parameters

**detailInfo** (STRING): Information regarding the circumstances in which the error was detected, for example, the API name and/or the piece of hardware that provided invalid data.

**invalidValue** (INHEX): The character in the string that was invalid, printed as a hexadecimal number.

## api.output.invalidSchema

### Severity

ERROR

### Description

This message occurs when an ONTAPI(tm) API generates output that violates the declared schema of the API. Most often this is due to a bug in the API handler.

### Corrective Action

The error is due to strict schema validation. You can turn off the output schema validation by using the following CLI commands: `$priv set advanced $registry set state.api.schema_output_validate.enable off`

### Syslog Message

Error in API output schema validation. API name: %s. Detailed error: %s.

### Parameters

**apiName** (STRING): Name of the ONTAPI API in which something invalid was detected.

**detailInfo** (STRING): Description of the error in the API output, such as an extra output value or a missing output value.

# app events

## app.cg events

### app.cg.snap.smbc.lookup.fail

#### Severity

ERROR

#### Description

This message occurs when a consistency group on either side of a SnapMirror® Business Continuity (SMBC) relationship cannot be resolved to the remote consistency group for correctly replicating consistency group Snapshot copies.

#### Corrective Action

Use the "volume snapshot show" command on the destination cluster to view the consistency group Snapshot copy and "volume snapshot restore" for restoring each individual volume instead of the consistency group REST interface.

#### Syslog Message

Failed to resolve Snapshot copy "%s" in consistency group "%s" on Vserver %s to the remote consistency group. Error: "%s".

#### Parameters

**snapshot\_name** (STRING): Name of the consistency group's Snapshot copy.

**consistency\_group** (STRING): Name of the consistency group.

**vserver** (STRING): Name of the Vserver.

**error** (STRING): The error that was encountered while resolving the local consistency group to its corresponding remote consistency group.

## app.log events

### app.log.alert

#### Severity

ALERT

#### Description

This event is issued when an application is in a condition that should be corrected immediately.

#### Corrective Action

(None).

#### Syslog Message

%s: %s %s: (%u) %s: %s

#### Parameters

**computerName** (STRING): Client Computer connected to the Filer.

**eventSource** (STRING): Client application that generated this event.

**appVersion** (STRING): Client application version.



**eventID** (INT): Application eventID.  
**category** (STRING): Event category.  
**subject** (STRING): Event description.

## app.log.emerg

### Severity

EMERGENCY

### Description

This message occurs when an application encounters a fatal error and requires immediate attention.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

%s: %s %s: (%u) %s: %s

### Parameters

**computerName** (STRING): Client Computer connected to the Filer.  
**eventSource** (STRING): Client application that generated this event.  
**appVersion** (STRING): Client application version.  
**eventID** (INT): Application eventID.  
**category** (STRING): Event category.  
**subject** (STRING): Event description.

## app.log.err

### Severity

ERROR

### Description

This event is issued when an application encounters an error condition.

### Corrective Action

(None).

### Syslog Message

%s: %s %s: (%u) %s: %s

### Parameters

**computerName** (STRING): Client Computer connected to the Filer.  
**eventSource** (STRING): Client application that generated this event.  
**appVersion** (STRING): Client application version.  
**eventID** (INT): Application eventID.  
**category** (STRING): Event category.  
**subject** (STRING): Event description.

## app.log.info

**Severity**

INFORMATIONAL

**Description**

This event is issued when an application's message is meant for informational purposes.

**Corrective Action**

(None).

**Syslog Message**

%s: %s %s: (%u) %s: %s

**Parameters**

**computerName** (STRING): Client Computer connected to the Filer.  
**eventSource** (STRING): Client application that generated this event.  
**appVersion** (STRING): Client application version.  
**eventID** (INT): Application eventID.  
**category** (STRING): Event category.  
**subject** (STRING): Event description.

**app.log.notice****Severity**

NOTICE

**Description**

This event is issued when an application is notifying of a certain event.

**Corrective Action**

(None).

**Syslog Message**

%s: %s %s: (%u) %s: %s

**Parameters**

**computerName** (STRING): Client Computer connected to the Filer.  
**eventSource** (STRING): Client application that generated this event.  
**appVersion** (STRING): Client application version.  
**eventID** (INT): Application eventID.  
**category** (STRING): Event category.  
**subject** (STRING): Event description.

**app.temp events****app.temp.delete.s3bucket.fail****Severity**

ERROR

## Description

This message occurs when a S3 Bucket cannot be deleted as a part of an application destroy operation.

## Corrective Action

Use the "object-store-server bucket delete" command to delete the S3 bucket manually. Retry the "application destroy" command on the specified application.

## Syslog Message

Failed to delete S3 Bucket %s of application %s on Vserver %s. Error %s.

## Parameters

**s3bucket** (STRING): Name of the S3 Bucket.

**application\_name** (STRING): Name of the application that contains this S3 bucket.

**vserver** (STRING): Name of the Vserver.

**error** (STRING): The error that the application destroy service encountered while destroying the S3 bucket.

## app.temp.destroy.flexcache.fail

### Severity

ERROR

### Description

This message occurs when a flexcache cannot be destroyed as a part of an application destroy operation.

### Corrective Action

Use the "volume offline" command to take the flexcache offline, and then use the "flexcache delete" command to delete the flexcache manually. Retry the "application destroy" command on the specified application.

### Syslog Message

Failed to destroy flexcache %s of application %s on Vserver %s. Error %s.

### Parameters

**flexcache** (STRING): Name of the flexcache.

**application\_name** (STRING): Name of the application that contains this volume.

**vserver** (STRING): Name of the Vserver.

**error** (STRING): The error that the application destroy service encountered while destroying the flexcache.

## app.temp.destroy.volume.fail

### Severity

ERROR

### Description

This message occurs when a volume cannot be destroyed as a part of an application destroy operation.

### Corrective Action

Use the "volume offline" command to take the volume offline, and then use the "volume delete" command to delete the volume manually. Retry the "application destroy" command on the specified application.

## Syslog Message

Failed to destroy volume %s of application %s on Vserver %s. Error %s.

## Parameters

**volume** (STRING): Name of the volume.

**application\_name** (STRING): Name of the application that contains this volume.

**vserver** (STRING): Name of the Vserver.

**error** (STRING): The error that the application destroy service encountered while destroying the volume.

## app.temp.mcc.sb.modify.abort

### Severity

ERROR

### Description

This message occurs on MetroCluster(tm) systems, as part of a failed "application modify" command. For this to happen, the primary site must be in the midst of modifying application provisioning when an unplanned MetroCluster switchover occurs. The secondary site must then successfully modify the same application. During MetroCluster switchback, the state of the original application provisioning job is checked before allowing the originally issued "application modify" command from the primary site to continue. If the state of the application has changed due to it being modified by the secondary site during switchover, the original "application modify" command is canceled during switchback. This is to prevent possible corruption to the application provisioning modifications from the secondary site.

### Corrective Action

Verify that the changes made by the MetroCluster secondary site are acceptable by using the "application show" command. If the changes are acceptable, no further action is required. Otherwise, rerun the "application modify" command as necessary.

## Syslog Message

Failed to modify application %s on Vserver %s.

## Parameters

**app\_name** (STRING): Name of the application being modified.

**vserver** (STRING): Name of the Vserver.

## app.temp.mcc.veto.fail

### Severity

ERROR

### Description

This message occurs when a metrocluster switchover or switchback command fails due to unfinished application jobs running on the node.

### Corrective Action

Run the "job show -category APS" command to first verify that there are no application provisioning jobs running. Then Retry the metrocluster switchover or switchback command.

## Syslog Message

A metrocluster switchover or switchback command was vetoed due to application provisioning jobs still in progress: %s.

**Parameters**

**error** (STRING): Error returned by Metrocluster application provisioning.

**app.temp.remove.app.fail****Severity**

ERROR

**Description**

This message occurs when an application cannot be removed from the persistent database during an "application destroy" operation.

**Corrective Action**

Retry the "application destroy" command. Contact NetApp technical support if the failure persists.

**Syslog Message**

Failed to remove application %s on Vserver %s from the persistent database. Error %s.

**Parameters**

**application\_name** (STRING): Name of the application.

**vserver\_name** (STRING): Name of the Vserver.

**error** (STRING): The error that the application destroy service encountered while removing the application from the persistent database.

# application events

## application.volume events

### application.volume.removed

#### Severity

ERROR

#### Description

This message occurs when an application is destroyed, but a member volume was not deleted after an error was encountered. This volume is no longer associated with the application and must be deleted manually.

#### Corrective Action

Delete the volume using the "volume delete" command.

#### Syslog Message

Failed to destroy member volume %s in Vserver %s during an application destroy operation. The volume must be deleted manually. The volume was previously associated with an application with UUID %s.

#### Parameters

**volumeName** (STRING): Name of the volume.

**vserver** (STRING): Vserver containing the volume.

**applicationUUID** (STRING): UUID of the application.

# arl events

## arl.abort events

### arl.abort.raid

#### Severity

ERROR

#### Description

This message occurs when a relocation of an aggregate is canceled due to activity in RAID.

#### Corrective Action

Reissue the "storage aggregate relocation start" command to initiate relocation of an aggregate to the specified destination node after the pending or in-progress operation on the aggregate is complete.

#### Syslog Message

Aggregate %s%s is %s; canceling 'storage aggregate relocation start' operation.

#### Parameters

**owner** (STRING): String indicating the owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.aggrdesttimeout events

### arl.aggrDestTimeout

#### Severity

NOTICE

#### Description

This message occurs during an aggregate relocation operation, when the source node does not receive notification from the destination node, within the specified time, that the aggregate is online.

#### Corrective Action

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online on the destination node.

#### Syslog Message

Node did not receive notification from the destination in %llu milliseconds that aggregate %s is online after relocation.

#### Parameters

**time** (LONGINT): Timeout, in milliseconds, for notification from the destination node that an aggregate is online after relocation.

**aggrName** (STRING): Name of the aggregate that was relocated.

# arl.aggronlinefailed events

## arl.aggrOnlineFailed

### Severity

EMERGENCY

### Description

This message occurs when the RAID or WAFL® subsystem fails to bring online the aggregate or its containing volumes after relocation.

### Corrective Action

Check the EMS logs for messages from the RAID and WAFL subsystems to determine the exact cause of the failure and the corrective action.

### Syslog Message

Bringing online of aggregate %s or its volumes failed after relocation with error %d.

### Parameters

**aggrName** (STRING): Name of the aggregate that was relocated.

**error** (INT): Error that occurred while bringing online the aggregate or its volumes after relocation.

# arl.aggronlinetimeout events

## arl.aggrOnlineTimeout

### Severity

ALERT

### Description

This message occurs when a node fails to bring online the aggregate and its containing volumes within the specified time after relocation.

### Corrective Action

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online.

### Syslog Message

Node did not bring online the aggregate %s and its volumes after relocation for %llu milliseconds.

### Parameters

**aggrUUID** (STRING): UUID of the aggregate that was relocated.

**time** (LONGINT): Timeout, in milliseconds, to bring online the aggregate and its volumes after relocation.

# arl.aggrstart events

## arl.aggrStart

### Severity

NOTICE



## Description

This message occurs when relocation of a storage failover aggregate starts.

## Corrective Action

(None).

## Syslog Message

Starting relocation of aggregate %s at time %llu.

## Parameters

**aggrName** (STRING): Name of the aggregate being relocated.

**time** (LONGINT): Time at which relocation of the aggregate started.

**aggrUUID** (STRING): UUID of the aggregate being relocated.

# arl.connecttoDestFailed events

## arl.connectToDestFailed

### Severity

ALERT

### Description

This message occurs when the local node cannot connect over the cluster network to the d-blade of the destination node while relocating an aggregate.

### Corrective Action

1. Verify that the destination node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local node and the destination node are both up.
3. Retry the aggregate relocation command.

### Syslog Message

Failed to connect with the destination node %s (d-blade ID: %s) over the cluster network while relocating aggregate %s.

### Parameters

**destName** (STRING): Name of the destination node.

**destDbladeid** (STRING): D-blade ID of the destination node.

**aggrName** (STRING): Name of the aggregate being relocated.

# arl.netra events

## arl.netra.ca.check.failed

### Severity

ERROR

### Description

This message occurs during the relocation of an aggregate, when the destination node cannot reach the object stores.

## Corrective Action

Verify that your intercluster LIF is online and functional by using the 'network interface show' command. Also, check network connectivity to the object store server by using the 'ping' command over the destination node intercluster LIF. Additionally, verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the 'aggregate object-store config show' command. It is possible to override this error by using the 'override-destination-checks' parameter of the relocation command. For more information or assistance, contact NetApp technical support.

## Syslog Message

Relocation of aggregate '%s' (uuid: %s) failed due to %s preventing object store access on the destination node.

## Parameters

**vol** (STRING): Name of the aggregate.  
**aggr\_uuid** (STRING): UUID of the aggregate.  
**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.netra.ha.dskChkFailed

### Severity

ALERT

### Description

This message occurs when relocation of an aggregate fails because the destination node cannot see all disks belonging to the aggregate.

### Corrective Action

If the destination is the source node's high-availability (HA) partner, identify the missing disks by using the "storage failover show -fields local-missing-disks, partner-missing-disks" command. If the destination is not the source node's HA partner, use the "storage disk show -aggregate" command to list the disks belonging to the aggregate, and then verify that those disks are visible from the destination by using the "storage disk show" command.

### Syslog Message

Relocation of aggregate %s failed from node %s because the destination node cannot see all disks belonging to the aggregate.

### Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.  
**node** (STRING): Name of the node that failed the relocation.

## arl.netra.lmgr.limit.failed

### Severity

ERROR

### Description

This message occurs during the relocation of an aggregate, when the destination node fails the lock count limit checks.

## Corrective Action

Ensure that the destination node supports the identified lock count limits, or select a different destination node and retry the operation.

## Syslog Message

Aggregate '%s': Lock limits check failed on destination node. %s

## Parameters

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.netra.lmgr.reserve.fail

### Severity

ERROR

### Description

This message occurs when an aggregate is being relocated, and the partner node fails to reserve lock resources needed for this aggregate.

### Corrective Action

Take the appropriate action based on the reason supplied in the logged message: - 'Partner reconstructing locks': Retry the relocation after reconstruction is complete on the partner node. You can check the reconstruction progress by using the 'debug locks reconstruction show' diagnostic privilege command. - 'Partner does not have sufficient lock resources': Free up lock resources on the partner and retry the relocation. You can free lock resources on the partner as client applications using locks are closed for volumes contained on that node. You can also free lock resources by moving volumes with locks from the partner node to a different node in the cluster. - 'Memory allocation failure': The system is running low on memory. Monitor memory statistics on the node and retry the relocation when more memory is available. - For other reason messages, or for further assistance, contact NetApp technical support.

## Syslog Message

Aggregate '%s': Lock reservation failed. Reason: %s.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.

**reason** (STRING): Reason the lock reservation failed.

## arl.netra.raid.failed

### Severity

ERROR

### Description

This message occurs during migration of an aggregate as part of aggregate relocation, when one of the aggregate checks fails on the destination node, which include duplicate aggregate name, duplicate Universally Unique Identifier (UUID), aggregate count, and capacity limits.

### Corrective Action

1. If checks fail due to aggregate count or capacity limits, then relocate this aggregate to an alternate destination node that has enough capacity to receive the aggregate.
2. If checks fail due to duplicate aggregate name or UUID, delete the duplicate aggregate created from previous "storage failover"

giveback" or "storage aggregate relocation" operations. Contact NetApp technical support for assistance with deletion of the unwanted duplicate aggregate.

### Syslog Message

Aggregate '%s' (UUID = %s, type = %s, home\_owner\_id = %llu, dr\_home\_owner\_id = %llu): RAID aggregate migration checks failed on destination node %s.

### Parameters

**vol** (STRING): Name of the aggregate.  
**aggregate\_uuid** (STRING): UUID of the aggregate.  
**raid\_type** (STRING): RAID type of the volume.  
**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.  
**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's disaster recovery (DR) home owner.  
**reason** (STRING): Activity that is preventing the aggregate relocation or giveback operation.

## arl.netra.wafl.mcc.veto

### Severity

ERROR

### Description

This message occurs during the relocation of an aggregate, when one or more online left-behind disaster recovery(DR) aggregates are found on the destination node and the "node-object-limit" option is set to off.

### Corrective Action

Check whether there exists any online left-behind DR aggregates on the destination node from a previous switchover operation. If such an aggregate exists, then perform the corrective actions specified in the earlier EMS messages for the left-behind aggregate to return it to its original owner, and then retry the operation. If you cannot perform the corrective action, then use the "override-destination-checks" option in the "relocation" command to force the relocation.

### Syslog Message

Aggregate '%s': One or more online left-behind DR aggregates were found on the destination node and the "node-object-limit" option is set to off.

### Parameters

**vol** (STRING): Name of the aggregate.

## arl.opaborted events

### arl.OpAborted

### Severity

ERROR

### Description

This message occurs when an aggregate relocation operation is aborted.

### Corrective Action

Refer to other generated error messages to identify the exact cause of failure and the corrective action.

## Syslog Message

Aggregate relocation operation aborted by node %s because %s.

## Parameters

**node** (STRING): Name of the node that aborted relocation.

**msg** (STRING): Error that prevented relocation.

# arl.opfinished events

## arl.OpFinished

### Severity

NOTICE

### Description

This message occurs when the aggregate relocation operation from the source node to the specified destination is finished. Use the "storage aggregate relocation show" command to check whether all aggregates were relocated successfully.

### Corrective Action

(None).

## Syslog Message

Aggregate relocation operation from the source node %s to the destination node %s finished in %llu milliseconds. 'override-vetoes' set to %s, and 'override-destination-checks' set to %s.

## Parameters

**srcName** (STRING): Name of the source node.

**destName** (STRING): Name of the destination node.

**time** (LONGINT): Time, in milliseconds, taken to complete the relocation operation.

**override\_vetoes** (STRING): Flag that indicates whether the system should override veto checks for a relocation operation. This flag corresponds to the "-override-vetoes" parameter of the "storage aggregate relocation start" command. When set to true, it might result in the relocation proceeding even if the node detects outstanding issues that would make aggregate relocation dangerous or disruptive.

**override\_dest\_checks** (STRING): Flag that indicates whether the system should override checks done on the destination node. This flag corresponds to the "-override-destination-checks" parameter of the "storage aggregate relocation start" command. When set to true, it can be used to force relocation of aggregates even if the destination has outstanding issues. Note that this could make the relocation dangerous or disruptive.

# arl.opstart events

## arl.OpStart

### Severity

NOTICE

### Description

This message occurs when the aggregate relocation operation starts. After this message is displayed, aggregates are relocated to the specified destination serially.

## Corrective Action

(None).

## Syslog Message

Starting relocation to destination %s at time %llu. 'override-vetoes' set to %s, and 'override-destination-checks' set to %s.

## Parameters

**destName** (STRING): Name of the destination node.

**time** (LONGINT): Time at which relocation started.

**override\_vetoes** (STRING): Flag that indicates whether the system should override veto checks for a relocation operation. This flag corresponds to the "-override-vetoes" parameter of the "storage aggregate relocation start" command. When set to true, it might result in the relocation proceeding even if the node detects outstanding issues that would make aggregate relocation dangerous or disruptive.

**override\_dest\_checks** (STRING): Flag that indicates whether the system should override checks done on the destination node. This flag corresponds to the "-override-destination-checks" parameter of the "storage aggregate relocation start" command. When set to true, it can be used to force relocation of aggregates even if the destination has outstanding issues. Note that this could make the relocation dangerous or disruptive.

# arl.postmigrnotok events

## arl.postMigrNotOk

### Severity

ERROR

### Description

This message occurs when the specified subsystem encounters an error during the post-migration phase of the aggregate relocation (ARL) operation.

### Corrective Action

Check the event log for a subsystem-specific reason as to why the subsystem failed the post-migration phase of the relocation operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by that message.

### Syslog Message

Relocation of aggregate '%s' encountered an error during post-migration processing. The error was encountered by the subsystem '%s'.

### Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the relocation operation.

**subsystem** (STRING): Name of the subsystem that failed the relocation operation.

# arl.precmt events

## arl.precmt.lmgr.resync.abrt

### Severity

NOTICE

## Description

This message occurs during aggregate relocation, when automatic lock resynchronization with the node's high-availability (HA) partner for files on this aggregate was aborted.

## Corrective Action

(None).

## Syslog Message

Automatic lock resynchronization was aborted for aggregate %s on forced relocation.

## Parameters

**aggr** (STRING): Name of the aggregate for which resynchronization was aborted.

## arl.precmt.lmgr.resyncing

### Severity

ERROR

### Description

This message occurs during aggregate relocation, when automatic lock resynchronization with the node's high-availability (HA) partner is in progress for files on this aggregate.

### Corrective Action

Retry the relocation after verifying that lock synchronization is complete by using the 'debug locks auto-resync-status -aggregate <aggregate>' diagnostic privilege command. If lock state disruption for all CA locks on the aggregate is acceptable, retry the aggregate relocation with the '-override-vetoes true' option.

### Syslog Message

Could not relocate aggregate %s because automatic lock resynchronization was in progress.

### Parameters

**aggr** (STRING): Name of the aggregate with locks not yet in sync.

## arl.precmt.lmgr.sync.failed

### Severity

ERROR

### Description

This message occurs when relocation of an aggregate has started but mirrored lock state cannot be created locally for the relocating aggregate. The relocate operation is aborted.

### Corrective Action

Take the appropriate action based on the reason supplied in the logged message: - 'Timed Out': Retry the aggregate relocation with the '-override-destination-checks' option. Overriding the check might result in a longer administrative outage period during which no administrative operations can be performed for volumes on this aggregate. - 'Fatal Error': Contact NetApp technical support.

### Syslog Message

Could not relocate the aggregate %s because creating locally mirrored lock state failed.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.  
**reason** (STRING): Reason the local mirrored lock state creation failed.

## arl.precmt.lmgr.sync.forced

### Severity

NOTICE

### Description

This message occurs during aggregate relocation, when local mirrored lock state creation is allowed to proceed without any time restrictions. This might result in a longer administrative outage period during which no administrative operations can be performed for volumes on this aggregate.

### Corrective Action

(None).

### Syslog Message

Local mirrored lock state creation for relocating aggregate %s is proceeding without time restrictions.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.

## arl.precmt.repl

### Severity

NOTICE

### Description

This message occurs when all replication transfers are aborted because of aggregate relocation.

### Corrective Action

(None).

### Syslog Message

Replication transfers aborted due to aggregate relocation.

## Parameters

(None).

## arl.precmt.vdom.lowMem

### Severity

ERROR

### Description

This message occurs when an aggregate relocation fails due to the unavailability of Data ONTAP® memory on the source node.



### **Corrective Action**

Free some memory and retry the aggregate relocation.

### **Syslog Message**

Relocation of aggregate %s failed due to unavailability of Data ONTAP® memory on the source node.

### **Parameters**

**aggr** (STRING): Name of the aggregate that was not relocated.

## **arl.precmt.wafl.volconversion**

### **Severity**

ERROR

### **Description**

This message occurs when an active volume conversion is running on one or more volumes on the source aggregate. Pre-commit is aborted.

### **Corrective Action**

When volume conversion is complete, re-run the command 'storage aggregate relocation start'.

### **Syslog Message**

Aggregate relocation/takeover is blocked as Volume Conversion is in progress on one of the volumes on aggregate %s.

### **Parameters**

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

## **arl.subsystemabort events**

### **arl.subsystemAbort**

### **Severity**

ERROR

### **Description**

This message occurs when the specified subsystem aborts relocation of the aggregate.

### **Corrective Action**

Check the syslog/EMS output for a subsystem-specific reason for aborting the relocation operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by the subsystem and then reissue the 'relocation' command. If you cannot perform the corrective action, then use the 'override-vetoes' option in the 'relocation' command to force the relocation.

### **Syslog Message**

The relocation operation of '%s' was aborted by '%s'.

## Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the relocation operation.

**subsystem** (STRING): Name of the subsystem that aborted the relocation operation.

# arl.veto events

## arl.veto.coredump.saveinprog

### Severity

INFORMATIONAL

### Description

This message occurs when the coredump subsystem vetoes the use of storage failover (SFO) during takeover because a coredump save is in progress; cluster failover (CFO) will be used instead. The coredump save must be complete before the system can perform a takeover.

### Corrective Action

(None).

### Syslog Message

Takeover use of SFO vetoed because a coredump save was in progress.

### Parameters

(None).

## arl.veto.coredump.unsavedcor

### Severity

ERROR

### Description

This message occurs when aggregate relocation (ARL) is vetoed because there is an unsaved kernel core file on one of the aggregate disks.

### Corrective Action

Use the "node coredump save" command to save any unsaved cores, or use the "-override-vetoes true" option when relocating aggregates.

### Syslog Message

Aggregate relocation vetoed due to an unsaved coredump %s on disk %s ("%s").

### Parameters

**corefile** (STRING): Name of the kernel core file.

**disk** (STRING): Name of the disk containing the kernel core file.

**panic** (STRING): Panic string of the unsaved kernel core file that caused the aggregate relocation (ARL) veto.

## arl.veto.kmgr.keysmissing

## Severity

ERROR

## Description

This message occurs when an aggregate relocation fails due to the unavailability of volume encryption keys for the encrypted volumes of the aggregate on the destination node.

## Corrective Action

Wait a few minutes, and then try the aggregate relocation again. If the problem persists, run the "security key-manager external restore" command for external key manager, or the "security key-manager onboard sync" command for Onboard Key Manager. The commands retrieve the encryption keys for the encrypted volumes of the aggregate from the key manager. Then try the aggregate relocation again.

## Syslog Message

Relocation of aggregate %s failed due to unavailability of volume encryption keys for the encrypted volumes of the aggregate on the destination node %s.

## Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.

**node** (STRING): Name of the destination node that is missing encryption keys.

## arl.veto.lmgr.CA.volmove

## Severity

ERROR

## Description

This message occurs when an aggregate relocation has started but cannot proceed because a volume move to or from this aggregate is in the cutover phase and the volume being moved has continuously available (CA) locks. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

## Corrective Action

Retry the relocation after volume move exits the cutover phase. To check for volume move state, run the 'volume move show' command.

## Syslog Message

Could not complete aggregate relocation because volume move of '%s%s%s%s' is in the cutover phase and this volume has CA locks.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## arl.veto.lmgr.nonCA.broken

**Severity**

NOTICE

**Description**

This message occurs when an aggregate relocation is forced that causes non-continuously available (non-CA) locks on the volume to be dropped. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

**Corrective Action**

(None).

**Syslog Message**

Dropped non-CA locks on volume %s%s%s%s due to forced aggregate relocation of aggregate %s.

**Parameters**

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

**aggrname** (STRING): Aggregate name.

**arl.veto.lmgr.nonCA.locks****Severity**

ERROR

**Description**

This message occurs when an aggregate relocation has started but cannot proceed because non-continuously available (non-CA) locks are present on the volume. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

**Corrective Action**

Based on how resilient they are to failures, applications must either gracefully close sessions over which non-CA locks are established or accept lock state disruption. To determine the open files that have these sessions established, run the 'vserver cifs session file show -hosting-aggregate "aggregate name" -continuously-available No' command. "aggregate name" is the aggregate for which relocation is being attempted. If lock state disruption for all existing non-CA locks is acceptable, retry the aggregate relocation by using the '-override-vetoes true' option.

**Syslog Message**

Could not complete aggregate relocation because of non-CA locks on volume %s%s%s%s aggregate %s.

**Parameters**

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

**aggrname** (STRING): Aggregate name.

## **arl.veto.lmgr.recons.left**

### **Severity**

ERROR

### **Description**

This message occurs when an aggregate relocation has started, but cannot proceed because reconstruction of client file locks from mirrored lock information is not yet complete.

### **Corrective Action**

Retry the relocation a few times after checking the status of reconstruction by using the 'debug locks reconstruction show' diagnostic privilege command. If the locks reconstruction operation is not completed, contact NetApp technical support.

### **Syslog Message**

Could not relocate aggregate '%s' because lock reconstruction on node %s is in progress.

### **Parameters**

**aggregate** (STRING): Name of the aggregate being relocated.

**node** (STRING): Name of the node reconstructing locks.

## **arl.veto.lmgr.syncing**

### **Severity**

ERROR

### **Description**

This message occurs when an aggregate relocation has started but cannot proceed because synchronization of client file locks ("locks sync") with the partner is not yet complete.

### **Corrective Action**

Retry the relocation after verifying that lock synchronization is complete, by using the 'debug locks lock-sync show' diagnostic privilege command.

### **Syslog Message**

Could not relocate aggregate '%s' because locks sync from node %s to its partner node %s is in progress.

### **Parameters**

**aggregate** (STRING): Name of the aggregate being relocated.

**node** (STRING): Name of the node performing locks sync with its partner.

**partner** (STRING): Name of the partner node with which locks sync is being performed.

## **arl.veto.repl**

### **Severity**

NOTICE

### **Description**

This message occurs when an aggregate relocation is vetoed by a critical replication transfer (for example,

volume move in the cutover phase).

### Corrective Action

Either abort the replication transfer that is preventing the aggregate relocation from being completed, retry the 'aggregate relocation start' command with the '-override-vetoes true' option, or wait for the transfer to finish.

### Syslog Message

Could not complete giveback because a replication transfer with UUID %s involving the aggregate %s is in progress.

### Parameters

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed the aggregate relocation.  
**aggr** (STRING): Name of the aggregate.

## arl.veto.snaprestore

### Severity

ERROR

### Description

This message occurs when the system cannot perform an aggregate relocation because a volume snaprestore operation is active. The relocation is aborted.

### Corrective Action

Retry aggregate relocation after the snaprestore operation is finished.

### Syslog Message

%s %s%s%s%s is running snaprestore. Canceling aggregate relocation.

### Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## arl.veto.volmove

### Severity

ERROR

### Description

This message occurs when an active volume move operation that cannot be automatically aborted prevents aggregate relocation from starting. When the volume move operation is completed, aggregate relocation can be retried using the 'storage aggregate relocation start' command.

### Corrective Action

When the volume move operation is complete, reissue the 'storage aggregate relocation start' command. To abort the volume move immediately, include the '-override-vetoes true' command option.

### **Syslog Message**

Unable to relocate aggregate while volume move for volume (DSID: %llu, NAME: %s) on aggregate %s is in progress.

### **Parameters**

**volume\_dsid** (LONGINT): Source/destination volume Data Set ID (DSID).

**vol\_name** (STRING): The name of the volume.

**aggr\_name** (STRING): The name of the containing aggregate.

## **arl.veto.wafl.volconversion**

### **Severity**

ERROR

### **Description**

This message occurs when an active volume conversion against one or more of the volumes on the source aggregate vetoed aggregate relocation.

### **Corrective Action**

When volume conversion is complete, re-run the command 'storage aggregate relocation start'.

### **Syslog Message**

Aggregate relocation is blocked as Volume Conversion is in progress on one of the volumes on aggregate %s.

### **Parameters**

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

# arw events

## arw.analytics events

### arw.analytics.ext.report

#### Severity

NOTICE

#### Description

This message occurs when anti-ransomware analytics generates or updates the "suspicious file extensions" report for a volume.

#### Corrective Action

(None).

#### Syslog Message

Anti-ransomware file extension analytics has found %s across high entropy data on the volume "%s" in Vserver "%s". Report location: %s.

#### Parameters

**fileExtension** (STRING): Previously unseen file extension or suspicious file extension

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**fileExtensionReport** (STRING): Link to report run after file extension analysis.

### arw.analytics.high.entropy

#### Severity

ERROR

#### Description

This message occurs when the number of high entropy data log messages (pertaining to ransomware detection and analysis) that were generated for a volume cross the predefined threshold.

#### Corrective Action

To correct this issue: 1) Restore data from the last Snapshot copy that were saved before high entropy data was detected. 2) Refer to the anti-ransomware documentation to learn how to implement ransomware protection and mitigation strategies.

#### Syslog Message

A large amount of high entropy data was found on volume "%s" in Vserver "%s".

#### Parameters

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

### arw.analytics.probability



**Severity**

ERROR

**Description**

This message occurs when an anti-ransomware attack probability has changed from "low" to "high" on a volume.

**Corrective Action**

To correct this issue: 1) Restore to the last safe Snapshot copy before attack probability became high. 2) Refer to the anti-ransomware documentation to diagnose further and take remedial measures.

**Syslog Message**

Anti-ransomware attack probability changed from low to high on volume "%s" in Vserver "%s".

**Parameters**

**volumeName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

**arw.analytics.report****Severity**

NOTICE

**Description**

This message occurs when an anti-ransomware analytics report is generated or updated for a volume.

**Corrective Action**

(None).

**Syslog Message**

Anti-ransomware analytics report has been generated for volume %s of Vserver %s. The report is available at %s.

**Parameters**

**volumeName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**report\_path** (STRING): Path to anti-ransomware report.

**arw.analytics.suspects****Severity**

ERROR

**Description**

This message occurs when a list of suspects generated by anti-ransomware analytics grows to a point where further investigation is needed.

**Corrective Action**

(None).

## Syslog Message

Anti-ransomware analytics on volume "%s" in Vserver "%s" has outstanding suspect files.

## Parameters

**volumeName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

# arw.auto events

## arw.auto.switch.enabled

### Severity

NOTICE

### Description

This message occurs when anti-ransomware has been automatically switched from learning mode to enabled after various conditions have been satisfied, such as learning period, file creation, file write, and file extension discovery activities.

### Corrective Action

(None).

## Syslog Message

Anti-ransomware has been automatically switched from learning mode to enabled on volume "%s" in SVM "%s".

## Parameters

**VolumeName** (STRING): Name of the volume.  
**VserverName** (STRING): Name of the SVM.  
**LearningPeriod** (INT): Number of days needed for learning.  
**IncomingWrite** (INT): Amount of write (in MB) received during Learning mode.  
**DurationWithoutNewExtension** (INT): Number of days without a new file extension being discovered.  
**FileCount** (INT): Number of new files created in a volume.  
**FileExtensionCount** (INT): Number of new file extensions discovered in a volume.

# arw.new events

## arw.new.file.extn.seen

### Severity

NOTICE

### Description

This message occurs when a new file extension is observed in anti-ransomware enabled volume. Its purpose is to promptly notify about the extension seen and enabling timely investigation.

### Corrective Action

(None).

## Syslog Message

A new file-extension "%s" is observed on volume "%s" (UUID: "%s") in SVM "%s" (UUID: "%s") at "%s".

## Parameters

**fileExtension** (STRING): Name of the File-Extension seen.  
**volumeName** (STRING): Name of the volume.  
**volumeUuid** (STRING): UUID of the Volume.  
**vserverName** (STRING): Name of the SVM.  
**vserverUuid** (STRING): UUID of the SVM.  
**time** (STRING): Time when new extension is seen.

# arw.snapshot events

## arw.snapshot.created

### Severity

NOTICE

### Description

This message occurs when a new ARP snapshot is created in anti-ransomware enabled volume. Additionally, it provides information about the reason behind the creation of the snapshot.

### Corrective Action

(None).

## Syslog Message

ARP snapshot created on volume "%s" (UUID: "%s") in SVM "%s" (UUID: "%s") at "%s". Reason: "%s".

## Parameters

**volumeName** (STRING): Name of the volume.  
**volumeUuid** (STRING): UUID of the Volume.  
**vserverName** (STRING): Name of the SVM.  
**vserverUuid** (STRING): UUID of the SVM.  
**time** (STRING): Time when new ARP snapshot is created.  
**reason** (STRING): Reason for snapshot creation.

# arw.volume events

## arw.volume.state

### Severity

NOTICE

### Description

This message occurs when the anti-ransomware state of a volume is changed.

### Corrective Action

(None).

## Syslog Message

Anti-ransomware state was changed to "%s" on volume "%s" (UUID: "%s") in Vserver "%s" (UUID: "%s").

## Parameters

**op** (STRING): Monitoring state (enabled, disabled, dry-run).  
**volumeName** (STRING): Name of the volume.  
**volumeUuid** (STRING): UUID of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**vserverUuid** (STRING): UUID of the Vserver.

# arw.vserver events

## arw.vserver.state

### Severity

NOTICE

### Description

This message occurs when the anti-ransomware state of a Vserver is changed.

### Corrective Action

(None).

## Syslog Message

Anti-ransomware was changed to "%s" on Vserver "%s" (UUID: "%s").

## Parameters

**op** (STRING): Anti-ransomware state (enabled, disabled, or dry-run).  
**vserverName** (STRING): Name of the Vserver.  
**vserverUuid** (STRING): UUID of the Vserver.

# asa events

## asa.personality events

### asa.personality.enable.end

**Severity**

NOTICE

**Description**

This message occurs when the cluster personality has been changed to an All SAN Array.

**Corrective Action**

(None).

**Syslog Message**

The cluster personality is now an All SAN Array.

**Parameters**

**job\_id** (LONGINT): ID of the conversion job.

### asa.personality.enable.fail

**Severity**

ERROR

**Description**

This message occurs when the job to change the cluster personality to an All SAN Array has failed.

**Corrective Action**

Attempt the failed command again. If the problem persists, contact NetApp technical support.

**Syslog Message**

The cluster personality was not changed to an All SAN Array. For more information, use the command "job show -id %llu" to see the job details.

**Parameters**

**job\_id** (LONGINT): ID of the conversion job.

### asa.personality.enable.start

**Severity**

NOTICE

**Description**

This message occurs when the cluster personality is being changed to All SAN Array.

**Corrective Action**

(None).

**Syslog Message**

The cluster personality is changing to All SAN Array.

**Parameters**

**job\_id** (LONGINT): ID of the conversion job.

# asup events

## asup.aods events

### asup.aods.response.overLimit

#### Severity

ERROR

#### Description

This message occurs when the size of the response from the AutoSupport OnDemand Server (AODS) goes over the AOD response buffer size limit in ONTAP.

#### Corrective Action

Contact Contact NetApp technical support. for assistance.

#### Syslog Message

AutoSupport OnDemand Server (AODS) response went over the response buffer size limit in ONTAP. Size limit: %u bytes.

#### Parameters

**SizeLimit** (INT): AOD response buffer size limit in bytes.

### asup.aods.response.timeOut

#### Severity

ERROR

#### Description

This message occurs when a response is not received from the AutoSupport OnDemand Server (AODS) before the system time limit.

#### Corrective Action

Check AOD Server connectivity using the "autosupport check show-details -check-type ondemand-server" command. If the issue persists, contact NetApp technical support.

#### Syslog Message

AutoSupport OnDemand Server (AODS) response was not received before the %u second time limit.

#### Parameters

**TimeLimit** (INT): System AODS response time limit in seconds.

## asup.config events

### asup.config.minimal.parse

#### Severity

INFORMATIONAL

## Description

This event is issued when we are unable to parse `/etc/asup_content.conf`

## Corrective Action

Either the `/etc/asup_content.conf` file was corrupt or it may have been accidentally modified, compare the file to the released version of the file and make the appropriate modifications.

## Syslog Message

AutoSupport could not parse section %s in `/etc/asup_content.conf`

## Parameters

**section** (STRING): The section that failed to parse.

## asup.config.minimal.unavailable

### Severity

ERROR

### Description

This event is issued when we are unable to access `/etc/asup_content.conf`

### Corrective Action

Either the `/etc/asup_content.conf` file does not exist or the permissions don't allow it to be accessed. Verify that the file is in the correct location and that the system has access to it.

### Syslog Message

Minimal AutoSupports unavailable. Could not read `/etc/asup_content.conf`

### Parameters

(None).

## asup.general events

### asup.general.drop

#### Severity

NOTICE

#### Description

This message occurs when the system cannot generate an AutoSupport message from a callhome EMS event.

#### Corrective Action

(None).

#### Syslog Message

Could not create AutoSupport message (sequence #: %u, ID: %s). %s.

#### Parameters

**sequence** (LONGINT): AutoSupport sequence number assigned to the callhome event.  
**trigger** (STRING): AutoSupport callhome EMS event ID.



**detail** (STRING): Additional details about why an AutoSupport message could not be created for the callhome EMS event.

**parameters** (STRING): A comma-delimited list of the parameter name and value pairs passed to the callhome EMS event.

## **asup.general.optout**

### **Severity**

NOTICE

### **Description**

This message occurs when an opt-out condition exists to notify you that you can turn off AutoSupport® reporting to NetApp, if desired. To disable AutoSupport reporting to NetApp, run the "system node autosupport modify -support disable" command within 24 hours of system initialization. Enabling AutoSupport can significantly speed problem determination and resolution, should a problem occur on your system.

### **Corrective Action**

(None).

### **Syslog Message**

AutoSupport will be enabled 24 hours after system initialization.

### **Parameters**

(None).

## **asup.general.queue.deleted**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when all entries have been deleted from the AutoSupport queue and the disk file spool. This should only occur when the autosupport.content option is changed from complete to minimal.

### **Corrective Action**

(None).

### **Syslog Message**

All AutoSupport message queue entries have been deleted.

### **Parameters**

(None).

## **asup.general.reminder**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when AutoSupport is not configured to send messages to NetApp.

## Corrective Action

(None).

## Syslog Message

AutoSupport is not configured to send messages to NetApp. Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system. More information about AutoSupport is available at <http://support.netapp.com/autosupport/>.

## Parameters

(None).

# asup.post events

## asup.post.drop

### Severity

ERROR

### Description

This message occurs when the AutoSupport® mechanism quits trying to post to the NetApp URL after having retried for the specified number of times.

### Corrective Action

Use the "system node autosupport check show-details" command to check for connectivity or certificate errors. If the connectivity check shows no errors and if the problem persists, contact NetApp technical support for further assistance.

### Syslog Message

AutoSupport message (%s) was not posted to NetApp. The system will drop the message.

### Parameters

**subject** (STRING): AutoSupport message title.

# asup.smtp events

## asup.smtp.drop

### Severity

ERROR

### Description

This message occurs when the AutoSupport® mechanism drops email messages after having retried for the specified times.

### Corrective Action

Use the "system node autosupport check show-details" command to check for connectivity errors with the configured mailhost. If the connectivity check shows no errors and if the problem persists, contact NetApp technical support for further assistance.

**Syslog Message**

AutoSupport mail (%s) was not sent. The system will drop the message.

**Parameters**

**subject** (STRING): AutoSupport message title.

## asup.throttle events

**asup.throttle.drop****Severity**

NOTICE

**Description**

This message occurs when an AutoSupport message is dropped because it is believed to be spam.

**Corrective Action**

To disable AutoSupport throttling, run the (privilege: diagnostic) "system node autosupport modify -throttle false -node <node name>" command.

**Syslog Message**

Too many AutoSupport messages in too short a time, throttling AutoSupport: %s

**Parameters**

**subject** (STRING): AutoSupport message title.

# asyncdelcli events

## asyncdelcli.message events

### asyncDelCli.message.finish

**Severity**

NOTICE

**Description**

This message occurs when an async-delete job has finished on the specified path.

**Corrective Action**

(None).

**Syslog Message**

Async-delete "%s" has finished on inode: %llu in volume (DSID: %llu, MSID: %llu).

**Parameters**

**Job\_Id** (STRING): Async-delete job ID.

**Inode** (LONGINT): Inode to be deleted.

**volume\_dsid** (LONGINT): DSID of the volume containing the directory.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

### asyncDelCli.message.start

**Severity**

NOTICE

**Description**

This message occurs when an async-delete job starts on the specified path.

**Corrective Action**

(None).

**Syslog Message**

Async-delete "%s" has started on inode: %llu in volume (DSID: %llu, MSID: %llu).

**Parameters**

**Job\_Id** (STRING): Async-delete job ID.

**Inode** (LONGINT): Inode to be deleted.

**volume\_dsid** (LONGINT): DSID of the volume containing the directory.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

## asyncdelcli.permscheck events

### asyncDelCli.permsCheck.fail

**Severity**

ALERT

**Description**

This message occurs when an async-delete job fails to delete file/directory due to permission issues.

**Corrective Action**

Verify that the user has permission to delete the inode.

**Syslog Message**

Async-delete %s encountered permission failures for user (%s) when deleting inode: %llu in volume (DSID: %llu, MSID: %llu).

**Parameters**

**Job\_Id** (STRING): Async-delete Job ID.

**User\_Id** (STRING): User Identifier.

**Path\_Inode** (LONGINT): Inode that has permission issues for deletion.

**volume\_dsid** (LONGINT): DSID of the volume containing the directory.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

# asyncdelete events

## asyncdelete.cancel events

### asyncDelete.cancel.start

#### Severity

NOTICE

#### Description

This message occurs when cancellation of async delete jobs is started on a node due to Vserver migrate operation. The administrator will need to restart the jobs after the migration is complete or aborted.

#### Corrective Action

The administrator will need to restart the cancelled async delete jobs after the migration is complete or aborted. Command to start an async-delete job: "volume file async-delete start -volume <volume-name> -vserver <vserver-name> -path <directory path>".

#### Syslog Message

Cancellation of async delete jobs was started on node "%s" because of Vserver migration. The administrator will need to restart the jobs after the migration is complete or aborted.

#### Parameters

**Node** (STRING): Node where async delete jobs are running.

## asyncdelete.message events

### asyncDelete.message.fail

#### Severity

ALERT

#### Description

This message occurs when an async delete job fails.

#### Corrective Action

Ensure that all nodes in the cluster are online and healthy, and that all volumes needed for the directory traversal are online. Then retry the operation.

#### Syslog Message

Async delete job on path %s of volume (MSID: %llu) failed: %s.

#### Parameters

**path** (STRING): Path to be deleted.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

**error** (STRING): Description of the error that occurred.

## **asyncDelete.message.start**

### **Severity**

NOTICE

### **Description**

This message occurs when an async delete job starts on the path specified.

### **Corrective Action**

(None).

### **Syslog Message**

Async delete job started on path %s of volume (MSID: %llu).

### **Parameters**

**path** (STRING): Path to be deleted.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

## **asyncDelete.message.success**

### **Severity**

NOTICE

### **Description**

This message occurs when an async delete job finishes successfully on the specified path.

### **Corrective Action**

(None).

### **Syslog Message**

Async delete job on path %s of volume (MSID: %llu) was completed. Number of files deleted: %llu, Number of directories deleted: %llu. Total number of bytes deleted: %llu.

### **Parameters**

**path** (STRING): Path to be deleted.

**volume\_msid** (LONGINT): MSID of the volume containing the directory.

**file\_count** (LONGINT): Total number of files deleted.

**dir\_count** (LONGINT): Total number of directories deleted.

**deleted\_size** (LONGINT): Total number of bytes deleted.

# ata events

## ata.excessive events

### ata.excessive.latency

**Severity**

NOTICE

**Description**

This message occurs when an I/O request to the ATA driver in ONTAP® Select takes longer than five seconds. This might be an indication of an overloaded or misconfigured back-end storage subsystem.

**Corrective Action**

If possible, reduce the I/O load on the system. Verify that the device is configured correctly and is functioning correctly.

**Syslog Message**

Excessive latency on ATA device %s

**Parameters**

`ata_latency_message` (STRING): ATA high latency message

## ata.timeout events

### ata.timeout.failed

**Deprecated**

Last used 9.2

**Severity**

NOTICE

**Description**

This message occurs when an I/O request to the ATA driver times out. This might be indicative of an overloaded back-end storage subsystem. The retries for this I/O operation have been exhausted; the I/O operation has failed and will not be retried.

**Corrective Action**

(None).

**Syslog Message**

ATA device timeout (%s).

**Parameters**

`ata_timeout` (STRING): Timeout message from the ATA driver.



## **ata.timeout.retry**

### **Deprecated**

This event is no longer used because of misleading text. See new EMS of EMS\_cam\_timeout\_retry. Last used 9.7.

### **Severity**

NOTICE

### **Description**

This message occurs when an I/O request to the ATA driver times out. This might be indicative of an overloaded back-end storage subsystem. The I/O operation will be retried; this message does not indicate failure.

### **Corrective Action**

(None).

### **Syslog Message**

ATA device timeout (%s).

### **Parameters**

**ata\_timeout** (STRING): Timeout message from the ATA driver.

# audit events

## audit.cap events

### audit.cap.staging.event.failure

**Severity**

NOTICE

**Description**

This message occurs when the system fails to generate an audit staging event as a part of the Central Access Policy (CAP) evaluation during access checks.

**Corrective Action**

(None).

**Syslog Message**

Failed to generate an audit staging event when evaluating CAPs for the Vserver "%s". Reason: "%s".

**Parameters**

**vserverUuidStr** (STRING): UUID of the Vserver.

**errorStr** (STRING): Error encountered.

## audit.log events

### audit.log.event

**Severity**

NOTICE

**Description**

This message occurs when there is an issue during message auditing. The system recovers by itself.

**Corrective Action**

(None).

**Syslog Message**

Auditing error (%s) occurred at %s.

**Parameters**

**error** (STRING): Error encountered during message auditing.

**time** (STRING): Time that this error occurred.

### audit.log.file.recovered

**Severity**

NOTICE

## Description

This message occurs when the system recovers a previously deleted log file as part of audit log policy.

## Corrective Action

(None).

## Syslog Message

Recovered previously deleted audit log file "%s".

## Parameters

**log\_name** (STRING): Log file name.

**recovered\_time** (STRING): File recovery timestamp.

## audit.log.queue.high

### Severity

NOTICE

### Description

This message occurs when the syslog queue reaches 75% of its maximum size.

### Corrective Action

Verify remote server network reachability and processing performance.

### Syslog Message

Log queue for destination %s, protocol %s has exceeded 75%% of capacity. Last audit log was sent successfully at %s, number of pending messages: %d.

### Parameters

**syslog\_dest** (STRING): Log forwarding destination for which messages are being dropped.

**protocol** (STRING): Protocol used for syslog destination.

**last\_sent** (STRING): Time that the audit log was last successfully sent to the syslog destination.

**pending\_msg** (INT): Number of pending messages in queue.

## audit.log.queue.overflow

### Severity

NOTICE

### Description

This message occurs when an internal queue overflow causes the audit logging mechanism to drop messages destined for a particular log forwarding destination.

### Corrective Action

Verify remote server network reachability and processing performance.

### Syslog Message

Log queue overflowed for destination %s, protocol %s. Current drop count: %ld.

## Parameters

**syslog\_dest** (STRING): Log forwarding destination for which messages are being dropped.

**protocol** (STRING): Protocol used for syslog destination.

**drop\_count** (LONGINT): Number of messages dropped since the last drop was reported.

## audit.log.queue.recovered

### Severity

NOTICE

### Description

This message occurs when the audit logging queue recovers after earlier message drops, and the current queue is less than 75% full.

### Corrective Action

(None).

### Syslog Message

Log queue recovered for destination %s, protocol %s. Issue start time: %s, end time: %s, total number of drops during this period: %ld.

## Parameters

**syslog\_dest** (STRING): Log forwarding destination for which the queue has recovered.

**protocol** (STRING): Protocol used for the syslog destination.

**issue\_start\_time** (STRING): Time that the queue overflowed.

**recovered\_time** (STRING): Time that the queue recovered.

**drop\_count** (LONGINT): Total number of drops since the last time the queue overflowed.

## audit.log.rotate.size.limit

### Severity

NOTICE

### Description

This message occurs when the system rotates a log file that has reached its size limit.

### Corrective Action

(None).

### Syslog Message

Log file %s has been rotated because size limit %ld was reached.

## Parameters

**log\_name** (STRING): Log file name.

**max\_log\_size** (LONGINT): Maximum log file size.

# auditlog events

## auditlog.change events

### auditlog.change.detected

#### Severity

ERROR

#### Description

This message occurs when the system detects that audit log files have been tampered with.

#### Corrective Action

Contact Contact NetApp technical support. for assistance.

#### Syslog Message

Audit log file "%s" was tampered with.

#### Parameters

**file\_name** (STRING): Name of the audit log file which was tampered with.

# auto events

## auto.cds events

### auto.cds.feature.disable.failed

#### Severity

INFORMATIONAL

#### Description

This message occurs during upgrade to Data ONTAP® 9.11.1 or later, when volume efficiency inactive data compression scan disablement failed on the existing compatible volumes.

#### Corrective Action

Manually try disabling the inline Volume Efficiency Inactive data compression on the volume using the "volume efficiency inactive-data-compression" command.

#### Syslog Message

Volume Inactive Data Compression disablement failed on Volume %s of Vserver %s.

#### Parameters

**volume** (STRING): Volume name.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for failure.

### auto.cds.feature.disable.success

#### Severity

INFORMATIONAL

#### Description

This message occurs during upgrade to Data ONTAP® 9.11.1 or later, when volume efficiency inactive data compression scan is disabled on the existing compatible volumes.

#### Corrective Action

(None).

#### Syslog Message

Volume Inactive Data Compression disabled on Volume %s of Vserver %s.

#### Parameters

**volume** (STRING): Volume name.

**vserver** (STRING): Vserver name.

## auto.policy events

### auto.policy.feature.enable

**Severity**

NOTICE

**Description**

This message occurs during upgrade to Data ONTAP® 9.3 or later, when we start enabling the "auto" efficiency policy on the existing compatible volumes on All-Flash Optimized nodes in the background.

**Corrective Action**

(None).

**Syslog Message**

Volume efficiency features will be enabled on Volumes of All-Flash Optimized nodes in the Cluster.

**Parameters**

(None).

**auto.policy.feature.enable.done****Severity**

NOTICE

**Description**

This message occurs during upgrade to Data ONTAP® 9.3 or later, when we have completed enabling the "auto" efficiency policy on the existing compatible volumes.

**Corrective Action**

(None).

**Syslog Message**

Volume efficiency features were enabled on Volumes of All-Flash Optimized nodes in the Cluster.

**Parameters**

(None).

**auto.policy.feature.enable.fail****Severity**

INFORMATIONAL

**Description**

This message occurs during upgrade to Data ONTAP® 9.3 or later, when we fail to enable the "auto" efficiency policy on the existing compatible volume on All-Flash Optimized nodes.

**Corrective Action**

Manually try enabling the inline Volume Efficiency features on the volume using the "volume efficiency modify" command.

**Syslog Message**

Could not enable Volume Efficiency features for Volume %s of vserver %s.%s

## Parameters

**volume** (STRING): Volume name.  
**vserver** (STRING): Vserver name.  
**reason** (STRING): Reason for failure.

## auto.policy.feature.enabled

### Severity

INFORMATIONAL

### Description

This message occurs during upgrade to Data ONTAP® 9.3 or later, when we enable the "auto" efficiency policy on the existing compatible volume on All-Flash Optimized nodes.

### Corrective Action

(None).

### Syslog Message

Volume Efficiency features enabled on Volume %s of Vserver %s.

## Parameters

**volume** (STRING): Volume name.  
**vserver** (STRING): Vserver name.

## auto.update events

### auto.update.action

#### Severity

NOTICE

#### Description

This message occurs when an action is applied to an automatic update request. This step requires input from the customer or a pre-configured default setting.

#### Corrective Action

(None).

#### Syslog Message

Automatic update: Applied the action "%s" to package: %s.

#### Parameters

**action** (STRING): Action selected for the automatic update request. Possible actions are "install", "schedule", or "dismiss".  
**package\_description** (STRING): Information about the package.  
**package\_id** (STRING): Unique identifier for the package.

### auto.update.eula.accepted



**Severity**

NOTICE

**Description**

This message occurs when the customer has accepted the End User License Agreement (EULA). This is required before the Automatic update feature is enabled.

**Corrective Action**

(None).

**Syslog Message**

Automatic update: End User License Agreement accepted by "%s".

**Parameters**

**user** (STRING): Name of the user account that accepted the EULA.

**auto.update.eula.reset****Severity**

NOTICE

**Description**

This message occurs when the Automatic update feature is disabled and the End User License Agreement (EULA) is no longer considered accepted.

**Corrective Action**

(None).

**Syslog Message**

Automatic update: Feature has been disabled by "%s".

**Parameters**

**user** (STRING): Name of the user account that disabled the EULA.

**auto.update.expired****Severity**

NOTICE

**Description**

This message occurs when an automatic update request has not been acted upon before the request expiration date.

**Corrective Action**

(None).

**Syslog Message**

Automatic update: Request expired on %s for package: %s.

**Parameters**

**expired\_date\_time** (STRING): Date and time when the request expired.

**package\_description** (STRING): Information about the package.

**package\_id** (STRING): Unique identifier for the package.

## auto.update.failure

### Severity

ERROR

### Description

This message occurs when an automatic update has failed to complete.

### Corrective Action

Fix the issues causing the failure and retry the automatic update. If the issue persists, contact NetApp technical support.

### Syslog Message

Automatic update: Failed update for package: %s. Error: %s.

### Parameters

**package\_description** (STRING): Information about the package.

**error\_msg** (STRING): String containing details regarding the cause of the failure.

**package\_id** (STRING): Unique identifier for the package.

**stage** (STRING): Stage at which the update failed.

## auto.update.received

### Severity

NOTICE

### Description

This message occurs when the automatic update information is received from the AutoSupport OnDemand server.

### Corrective Action

(None).

### Syslog Message

Automatic update: Received update for package: %s.

### Parameters

**package\_description** (STRING): Information about the package.

**package\_id** (STRING): Unique identifier for the package.

## auto.update.retracted

### Severity

NOTICE

### Description

This message occurs when an automatic update request has been superseded, retracted or otherwise considered unnecessary.

**Corrective Action**

(None).

**Syslog Message**

Automatic update: Request retracted for package: %s.

**Parameters**

**package\_description** (STRING): Information about the package.

**package\_id** (STRING): Unique identifier for the package.

**auto.update.success****Severity**

NOTICE

**Description**

This message occurs when an automatic update has completed successfully.

**Corrective Action**

(None).

**Syslog Message**

Automatic update: Successfully updated package: %s.

**Parameters**

**package\_description** (STRING): Information about the package.

**package\_id** (STRING): Unique identifier for the package.

# aws events

## aws.dev events

### aws.dev.diskFailure

#### Severity

ERROR

#### Description

This message occurs when ONTAP® is running as an instance in Amazon Web Services (AWS), and a data disk failure occurs because of an I/O hang. ONTAP marks the disk as broken and marks the data aggregate as partial or degraded.

#### Corrective Action

Detach the failed device from the AWS instance.

#### Syslog Message

Disk failure detected after I/O to device %s hung for %d seconds.

#### Parameters

**dev** (STRING): Name of the device.

**timeout** (INT): Number of seconds the disk is hung.

### aws.dev.diskNotUnique

#### Severity

ERROR

#### Description

This message occurs when Data ONTAP® is running as an instance in Amazon Web Services (AWS), and the volume being added has a UUID that is already used by the attached volumes. This usually happens when the user creates a Snapshot® copy from a volume, creates the volume from this copy and then tries to attach the volume to the instance.

#### Corrective Action

Detach the being added duplicate UUID device from the AWS instance.

#### Syslog Message

The volume #%d(%s) being added has a duplicate UUID which is already used by the attached volume #%d(%s).

#### Parameters

**curr\_unit** (INT): Unit number of the device which is being attached.

**curr\_devname** (STRING): Name of the device which is being attached.

**orig\_unit** (INT): Unit number of the original device which is being duplicated.

**orig\_devname** (STRING): Name of the original device which is being duplicated.

## aws.dev.limitExceeded

### Severity

ERROR

### Description

This message occurs when Data ONTAP® is running as an instance in Amazon Web Services (AWS), and the volume being added is attached to an unsupported device name. Data ONTAP supports attaching a maximum of 51 Elastic Block Storage (EBS) volumes to a single instance. Supported device names are xvdf-xvdz, xvdaa-xvdaz, and xvdba-xvdbd.

### Corrective Action

Detach the unsupported device from the AWS instance.

### Syslog Message

The volume being added is attached to an unsupported device name `##d(%)s`. Data ONTAP supports attaching a maximum of 51 device volumes to a single instance. Supported device names are xvdf-xvdz, xvdaa-xvdaz, and xvdba-xvdbd.

### Parameters

**unit** (INT): Device unit number.  
**devname** (STRING): Device name.

## aws.instance events

### aws.instance.metadata.fail

### Severity

ERROR

### Description

This message occurs when instance metadata retrieval fails on a hypervisor hosting an Amazon Web Services (AWS) ONTAP instance. The Instance Metadata Service (IMDS) is not reachable at this time. Failure to retrieve the instance metadata indicates issues with the hypervisor.

### Corrective Action

This error indicates an issue with the AWS environment. Contact AWS technical support for further assistance. Stopping and starting the instance can also resolve the issue because the instance will typically be relaunched on a different host.

### Syslog Message

IMDS%s request fails to retrieve AWS instance metadata: %s.

### Parameters

**IMDS\_Version** (STRING): Metadata request version that generates this event.  
**metadataType** (STRING): Metadata request type that generates this event.

## aws.nic events

## aws.nic.noTxIntr

### Severity

ALERT

### Description

This message occurs when Data ONTAP® is running in Amazon Web Services (AWS) and a network interface stops processing packets at the beginning of system boot.

### Corrective Action

If the issue persists, reboot the node. If reboot does not help, contact NetApp technical support for troubleshooting assistance.

### Syslog Message

No TX interrupt on NIC '%s' for %u seconds.

### Parameters

**nicName** (STRING): Network interface or port name.  
**secElapsed** (INT): Seconds elapsed since the last send interrupt.  
**sigCnt** (INT): Number of signals sent to the back-end host OS.  
**reqPending** (INT): Number of outstanding send requests.  
**reqTotal** (INT): Number of total send requests.  
**respPending** (INT): Number of outstanding responses.  
**respTotal** (INT): Number of total responses.  
**evtBackend** (INT): Number of back-end events.  
**evtFrontend** (INT): Number of front-end events.

# backup events

## backup.abort events

### backup.abort.ROOLR

**Severity**

NOTICE

**Description**

This event is generated when backup event like dump has been aborted and currently ROOLR (Read Only On Line Replica) is running on the Filer. This event might have been aborted because of user intervention or device like bad tape has encountered backup failure.

**Corrective Action**

(None).

**Syslog Message**

The abort event, %s, is just notified.

**Parameters**

**errmsg** (STRING): Error Description

### backup.abort.STSP

**Severity**

NOTICE

**Description**

This event is generated when a backup started by VERITAS NBU STSP Protocol has been aborted. This event might have been aborted because of user intervention or device like bad tape has encountered backup failure.

**Corrective Action**

(None).

**Syslog Message**

The abort event, %s, is just notified.

**Parameters**

**errmsg** (STRING): Error Description

## backup.fail events

### backup.fail.giveback

**Severity**

ERROR

## Description

This message occurs during a giveback operation, when there are active dump or restore processes running on cfo aggregates of the partner node.

## Corrective Action

Wait for backup operations like dump or restore running on cfo aggregates of the partner node to finish before retrying the giveback operation. To kill the active backup operations and proceed, retry the giveback operation with the `-override-vetoes` option set to true.

## Syslog Message

Backup/restore services: There are active backup/restore sessions on the partner.

## Parameters

(None).

## backup.fail.logopen

### Severity

ERROR

### Description

This event is generated when we tried to open up the log file and it failed while opening the file

### Corrective Action

Make sure that the log file exists with proper permission and the file is not corrupted.

### Syslog Message

Failed to open backup log file: %s

### Parameters

`log_file_open_err` (STRING): The error encountered during log file open

## backup.fail.logwrite

### Severity

ERROR

### Description

This event is generated when we tried to write a log event to the log file and it failed while writing it to the file.

### Corrective Action

Make sure that the log file exists with proper permission and the file is not corrupted.

### Syslog Message

Failed to write to backup log file: %s

### Parameters

`log_file_write_err` (STRING): The error encountered during log file write



# backup.flip events

## backup.flip.inodemap

### Severity

INFORMATIONAL

### Description

This message is generated while flipping inodemap.

### Corrective Action

(None).

### Syslog Message

Took %u seconds to flip inodemap.

### Parameters

**time** (INT): Time in seconds taken to flip inodemap.

# bcomd events

## bcomd.fcNvmelifadded events

### bcomd.fcNvmeLifAdded

#### Severity

NOTICE

#### Description

This message occurs when an FC NVMe (Nonvolatile Memory Express) logical interface (LIF) is added.

#### Corrective Action

(None).

#### Syslog Message

LIF %s (on Vserver %s), WWPN %s, was added to node %s, port %s.

#### Parameters

**lifname** (STRING): Name of the LIF that was added.  
**vserver** (STRING): Name of the Vserver on which the LIF resides.  
**wwpn** (STRING): WWPN that was added.  
**node** (STRING): Name of the node on which the LIF is hosted.  
**port** (STRING): Name of the port on which the LIF is hosted.

## bcomd.fcNvmelifmovefail events

### bcomd.fcNvmeLifMoveFail

#### Severity

EMERGENCY

#### Description

This message occurs when a LIF move failed and the recovery action to restore the LIF has also failed. Namespaces are no longer accessible through this LIF. The LIF must be recovered manually.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Moving FC-NVMe LIF %s of vserver %s from %s:%s to %s:%s failed. Manual recovery required.

#### Parameters

**lif\_name** (STRING): LIF name.  
**vserver\_uuid** (STRING): Vserver Uuid.  
**src\_node** (STRING): LIF was moving from this node.  
**src\_port** (STRING): LIF was moving from this port.  
**dest\_node** (STRING): LIF was moving to this node.  
**dest\_port** (STRING): LIF was moving to this port.

## bcomd.fcNvmelifremoved events

### bcomd.fcNvmeLifRemoved

#### Severity

NOTICE

#### Description

This message occurs when an FC NVMe (Nonvolatile Memory Express) logical interface (LIF) is removed.

#### Corrective Action

(None).

#### Syslog Message

LIF %s (on Vserver %s), WWPN %s, was removed from node %s, port %s.

#### Parameters

**lifname** (STRING): Name of the LIF that was removed.

**vserver** (STRING): Name of the Vserver on which the LIF resided.

**wwpn** (STRING): WWPN that was removed.

**node** (STRING): Name of the node on which the LIF had been hosted.

**port** (STRING): Name of the port on which the LIF had been hosted.

## bcomd.fcplifadded events

### bcomd.fcpLifAdded

#### Severity

NOTICE

#### Description

This message occurs when a Fibre Channel logical interface (LIF) is added.

#### Corrective Action

(None).

#### Syslog Message

LIF %s (on Vserver %s), WWPN %s, was added to node %s, port %s.

#### Parameters

**lifname** (STRING): Name of the LIF that was added.

**vserver** (STRING): Name of the Vserver on which the LIF resides.

**wwpn** (STRING): WWPN that was added.

**node** (STRING): Name of the node on which the LIF is hosted.

**port** (STRING): Name of the port on which the LIF is hosted.

## bcomd.fcplifremoved events

## bcomd.fcpLifRemoved

### Severity

NOTICE

### Description

This message occurs when a Fibre Channel logical interface (LIF) is removed.

### Corrective Action

(None).

### Syslog Message

LIF %s (on Vserver %s), WWPN %s, was removed from node %s, port %s.

### Parameters

**lifname** (STRING): Name of the LIF that was removed.

**vserver** (STRING): Name of the Vserver on which the LIF resided.

**wwpn** (STRING): WWPN that was removed.

**node** (STRING): Name of the node on which the LIF had been hosted.

**port** (STRING): Name of the port on which the LIF had been hosted.

## bcomd.igroup events

### bcomd.igroup.auto.delete.fail

### Severity

ERROR

### Description

This message occurs when an initiator group marked with the delete-on-unmap flag is not automatically deleted when it is no longer part of a LUN mapping relationship.

### Corrective Action

Use the "lun igroup delete" command to delete the initiator group manually.

### Syslog Message

Failed to delete initiator group %s on Vserver %s: %s

### Parameters

**igroup** (STRING): Name of the initiator group.

**vserver** (STRING): Name of the Vserver.

**error** (STRING): The error encountered when deleting the initiator group.

## bcomd.lun events

### bcomd.lun.update.lost

### Severity

ERROR

## Description

This message occurs when the SAN Kernel Agent is unable to process a LUN metadata update resulting from a Snapshot® copy restore operation or other LUN event.

## Corrective Action

Resolve the error and then reload the indicated Vserver with the command "debug san vserver reload".

## Syslog Message

A metadata change notification for a LUN in VServer %s on volume %llu has been lost due to error: %s

## Parameters

**vserver** (STRING): Name of the Vserver.

**volume** (LONGINT): MSID of the volume.

**error** (STRING): Error preventing the update from being processed.

# bcomd.vdisk events

## bcomd.vdisk.repl.failed

### Severity

ERROR

### Description

This message occurs after a replication failure of a LUN or an NVMe namespace configuration in a MetroCluster™ configuration.

### Corrective Action

Use the "metrocluster vserver resync" command to resynchronize the Vserver configuration for the affected Vserver.

### Syslog Message

Replication of a LUN or an NVMe namespace configuration failed on Vserver "%s". Reason: %s.

### Parameters

**vserver** (STRING): Name of the Vserver.

**error** (STRING): Error encountered while replicating the configuration.

## bcomd.vdisk.transition

### Severity

NOTICE

### Description

This message occurs when a user attempts to convert a LUN into an NVMe namespace or vice-versa.

### Corrective Action

(None).

### Syslog Message

Conversion for "%s", Vserver "%s" to type "%s" completed with status "%s".

## Parameters

**source\_vdisk\_path** (STRING): Path to the LUN or NVMe namespace to be converted.  
**vserver** (STRING): Name of the Vserver.  
**destination\_vdisk\_type** (STRING): Destination vdisk type after conversion.  
**status** (STRING): Status of conversion operation.

# bcomd.volume events

## bcomd.volume.update.lost

### Severity

ERROR

### Description

This message occurs when the SAN Kernel Agent is unable to process a volume LUN inventory update resulting from a Snapshot® copy restore operation.

### Corrective Action

Resolve the error and then delete stale maps in the volume with the command "lun mapping delete-stale".

### Syslog Message

A volume LUN inventory change notification for Vserver %s volume %llu has been lost due to error: %s

## Parameters

**vserver** (STRING): Name of the Vserver.  
**volume** (LONGINT): MSID of the volume.  
**error** (STRING): Error preventing the update from being processed.

# bdfu events

## bdfu.selected events

### bdfu.selected

#### Severity

INFORMATIONAL

#### Description

This message occurs when background disk firmware update selects a disk for firmware download. The disk contents will be copied to a suitable spare, the disk will be removed from the raid group and then firmware will be downloaded to it. After the firmware download is complete, the contents of the disk will be restored.

#### Corrective Action

(None).

#### Syslog Message

Disk %s [%s %s %s] S/N [%s] selected for background disk firmware update.

#### Parameters

**diskName** (STRING): Name of the disk.

**vendorName** (STRING): Vendor name of the disk.

**productId** (STRING): Product ID of the disk.

**fwVersion** (STRING): Firmware version of the disk.

**serialno** (STRING): Serial number of the disk.

# block events

## block.giveback events

### block.giveback.during.aufo

#### Severity

ALERT

#### Description

This message occurs during a storage failover (SFO) aggregate giveback, when a SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation is in progress.

#### Corrective Action

Wait for the operation to finish. If giveback is required immediately, use the "override-veto true" parameter with the giveback command, or retry the operation with the "override-vetoes true" parameter. Note: Using this option might cause an I/O disruption.

#### Syslog Message

Storage Failover giveback cannot complete for aggregate "%s" because an SMBC automatic unplanned failover is in progress.

#### Parameters

**aggrname** (STRING): Name of the aggregate.



# boot events

## boot.bootmenu events

### boot.bootmenu.issue

#### Severity

EMERGENCY

#### Description

This message occurs when an error is detected while booting into, or using, the system boot menu.

#### Corrective Action

Retry the boot. If the problem persists, contact NetApp technical support for assistance in replacing the boot device media.

#### Syslog Message

Failed to boot the system. Error: %s

#### Parameters

`err_str` (STRING): Error string.

## boot.shutdown events

### boot.shutdown.issue

#### Severity

ALERT

#### Description

This message occurs when an error is detected while shutting down the system.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Failed to shutdown the system. Error: %s

#### Parameters

`err_str` (STRING): Error string.

## boot.varfs events

### boot.varfs.backup.issue

#### Severity

EMERGENCY

**Description**

This message occurs when an error is detected while trying to back up the /var file system to the boot media or node root volume (mroot).

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Backup of the /var file system failed (%s)

**Parameters**

**err\_str** (STRING): Error string.

**boot.varfs.backup.success****Severity**

NOTICE

**Description**

This message occurs when a subsequent backup is successful following a failure.

**Corrective Action**

(None).

**Syslog Message**

Backup of the /var file system is successful after previous failures.

**Parameters**

(None).

**boot.varfs.repair.partial****Severity**

ERROR

**Description**

This message occurs when the system is unable to restore all missing files detected by a scan of the /var file system during boot.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Unable to restore all missing files in /var file system.

**Parameters**

(None).

**boot.varfs.repair.success**

**Severity**

NOTICE

**Description**

This message occurs when missing files are restored during a scan of the /var file system during boot.

**Corrective Action**

(None).

**Syslog Message**

All missing files in /var are restored.

**Parameters**

(None).

**boot.varfs.restore.issue****Severity**

ALERT

**Description**

This message occurs when an error is detected while restoring the /var file system from a shutdown.

**Corrective Action**

The file system is automatically restored from the boot media. If the problem persists, contact NetApp technical support.

**Syslog Message**

Failed to restore the /var file system. Error: %s

**Parameters**

**err\_str** (STRING): Error string.

**boot.varfs.restore.success****Severity**

NOTICE

**Description**

This message occurs when, following a failure, a subsequent restore of the /var file system is successful.

**Corrective Action**

(None).

**Syslog Message**

Successfully restored the /var file system.

**Parameters**

(None).

# boot.verify events

## boot.verify.kmods.disabled

### Severity

ALERT

### Description

This message occurs when the verification of digitally signed kernel modules has been disabled.

### Corrective Action

Contact NetApp technical support immediately.

### Syslog Message

The verification of digitally signed kernel modules has been disabled.

### Parameters

(None).

## boot.verify.kmods.success

### Severity

NOTICE

### Description

This message occurs when the verification of digitally signed kernel modules succeeds.

### Corrective Action

(None).

### Syslog Message

The verification of digitally signed kernel modules has completed.

### Parameters

(None).

# bootfs events

## bootfs.env events

### bootfs.env.issue

#### Severity

ERROR

#### Description

This message occurs when an error is detected while trying to access the environment file stored on the boot device. Environment variables on the system cannot be changed or saved.

#### Corrective Action

Reformat or replace the boot media. If the problem persists, contact NetApp technical support.

#### Syslog Message

Environmental variables cannot be accessed or saved on the boot device (%s).

#### Parameters

**err\_str** (STRING): Error string.

# bridge events

## bridge.discovered events

### bridge.discovered

#### Severity

INFORMATIONAL

#### Description

This message occurs when the system discovers an FC-to-SAS bridge.

#### Corrective Action

(None).

#### Syslog Message

FC-to-SAS bridge %s [%s %s %s] S/N [%s] was discovered.

#### Parameters

**bridgeName** (STRING): Name of the bridge.  
**bridgeVendor** (STRING): Vendor name of the bridge.  
**bridgeModel** (STRING): Model name of the bridge.  
**bridgeRevision** (STRING): Firmware revision of the bridge.  
**bridgeSerialNum** (STRING): Serial number of the bridge.

## bridge.duplicate events

### bridge.duplicate.ip.address

#### Severity

ALERT

#### Description

This message occurs when the system discovers an ATTO FibreBridge using the same IP address as some other device in the network.

#### Corrective Action

Using the IP and MAC addresses from the alert, attempt to locate the other device in the network, and change one of the IP addresses. If unsuccessful, contact NetApp support for assistance with this issue.

#### Syslog Message

FC-to-SAS bridge %s [%s] %s %s %s.

#### Parameters

**bridgeName** (STRING): Name of the bridge that detected the duplicate IP address.  
**bridgeSerialNum** (STRING): Serial number of the bridge that detected the duplicate address.  
**duplicateIPAddress** (STRING): Duplicate IP address detection message from the bridge log.  
**duplicateMAC** (STRING): Duplicate MAC address detection message from the bridge log.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## bridge.duplicate.wwn

### Severity

ALERT

### Description

This message occurs when the system discovers two FC-to-SAS bridges reporting the same WWN.

### Corrective Action

Contact NetApp support for assistance with this issue.

### Syslog Message

FC-to-SAS bridge %s [%s] and bridge %s [%s] are presenting a duplicate WWN [%s] %s.

### Parameters

**bridgeName1** (STRING): Name of the first bridge.

**bridgeSerialNum1** (STRING): Serial number of the first bridge.

**bridgeName2** (STRING): Name of the second bridge.

**bridgeSerialNum2** (STRING): Serial number of the secondbridge.

**bridgeWWN** (STRING): World wide name of the bridges.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

## bridge.removed events

### bridge.removed

### Severity

INFORMATIONAL

### Description

This message occurs when an FC-to-SAS bridge is removed from the system.

### Corrective Action

(None).

### Syslog Message

FC-to-SAS bridge %s [%s %s %s] S/N [%s] was removed.

### Parameters

**bridgeName** (STRING): Name of the bridge.

**bridgeVendor** (STRING): Vendor name of the bridge.

**bridgeModel** (STRING): Model name of the bridge.

**bridgeRevision** (STRING): Firmware revision of the bridge.

**bridgeSerialNum** (STRING): Serial number of the bridge.

## bridge.sas events

### bridge.SAS.invalidTransition

**Severity**

ERROR

**Description**

This message occurs when the SAS disk shelves behind an FC-to-SAS bridge are not cabled correctly, indicating an invalid SAS2 to SAS3 transition.

**Corrective Action**

Follow the defined Fabric-attached or stretch MetroCluster cabling guide.

**Syslog Message**

An invalid SAS2 to SAS3 transition exists on FC-to-SAS bridge %s S/N [%s]: %s

**Parameters**

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeError** (STRING): Error string from the FC-to-SAS bridge.

## bridge.timestamp events

### bridge.timestamp.noSync

**Severity**

NOTICE

**Description**

This message occurs when the max number of retries is exceeded when attempting to apply SCSI SET TIMESTAMP command on one or more attached FC-to-SAS bridges. A timestamp that is not in sync does not affect the functional operation of the bridge. The SET TIMESTAMP operation is only used to synchronize timestamps between ONTAP® and the bridge to assist with correlating log events.

**Corrective Action**

A timestamp operation will always retry in 24 hours. If this notice persists, contact NetApp technical support.

**Syslog Message**

ONTAP timestamp sync did not complete on FC-to-SAS bridge with S/N: %s WWN: %s.

**Parameters**

**bridgeSerialNum** (STRING): Serial number of the bridge.

**bridgeWWN** (STRING): World Wide Name of the bridge.



# bsdsocket events

## bsdsocket.queueLimit events

### bsdsocket.queueLimit.reached

#### Severity

ERROR

#### Description

This message occurs when the number of connection requests for a socket reaches the maximum. These connection requests include partially completed connections and connections that are ready to be accepted.

#### Corrective Action

Network misconfigurations can cause connectivity issues such as asymmetric routing, which causes loss of ACKs. Check the network for any such misconfigurations. A Denial of Service attack and an overloaded system can also be reasons for such an event. Contact NetApp technical support.

#### Syslog Message

Number of connection requests for the socket bound to port %d having the IP address %s has reached the limit of %d connections. The socket uses %s protocol. This event has occurred %llu times since the last warning.

#### Parameters

**port** (INT): Local port bound to the socket.

**ipAddress** (STRING): IP address of the socket.

**limit** (INT): Maximum number of connection requests per socket.

**protocol** (STRING): Protocol used by the socket.

**count** (LONGINT): Number of occurrences for this event since the last warning.

# bucket events

## bucket.soft events

### bucket.soft.limit.exceeded

#### Severity

NOTICE

#### Description

This message occurs when a bucket exceeds the soft quota limit (block).

#### Corrective Action

(None).

#### Syslog Message

Bucket "%s" on volume "%s%s%s" has exceeded the soft block limit and is at %s capacity.

#### Parameters

**whom** (STRING): Entity for which this event is generated.

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**capacity** (STRING): Capacity percentage.

### bucket.soft.limit.normal

#### Severity

NOTICE

#### Description

This message occurs when a bucket's soft quota limit (block) returns to normal.

#### Corrective Action

(None).

#### Syslog Message

Bucket "%s" on volume "%s%s%s" is below the soft block limit and is at %s capacity.

#### Parameters

**whom** (STRING): Entity for which this event is generated.

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**capacity** (STRING): Capacity percentage.

# cache events

## cache.ttl events

### cache.ttl.settings.change

**Severity**

ERROR

**Description**

This message occurs on the last node during the upgrade to ONTAP® 9.3, if the default TTL values of any of the SecD or NFS export caches are changed.

**Corrective Action**

Use the "vserver services name-service cache" commands to configure TTL values after the system is fully upgraded to ONTAP 9.3.

**Syslog Message**

One or more default cache TTL values for SecD and NFS export caches have been changed.

**Parameters**

(None).

# cachemirror events

## cachemirror.offline events

### cacheMirror.offline

#### Severity

ERROR

#### Description

This message occurs when Cache Mirror encounters an error and has to be taken offline. The system functions normally except that takeover and giveback operation times will be negatively affected.

#### Corrective Action

Reboot the node by performing a planned takeover and giveback. Contact NetApp technical support if the problem persists.

#### Syslog Message

Cache Mirror was taken offline because of an unrecoverable error. %s.

#### Parameters

**error\_reason** (STRING): String providing information about why Cache Mirror was taken offline.

# callhome events

## callhome.aggr events

### callhome.aggr.meta.repair

#### Severity

ERROR

#### Description

This message occurs after WAFL® Iron (wafliron) detects and repairs corruption in an aggregate metafile. Wafliron is a detection and repair process. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Wafliron might need to be run on other volumes in aggregate or complete aggregate. Check for wafl.iron.cont.corrupt messages in the node's event log to find other corrupt volumes. Run wafliron on all corrupt volumes. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for WAFLIRON AGGREGATE METAFILE REPAIRED

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.aggr.offline

#### Severity

ALERT

#### Description

This message occurs when an aggregate is inconsistent and is taken offline. The inconsistency might stem from one or more sources and might be related to previous operations or to current operations. Disabling access is a protective measure to ensure that further inconsistencies are avoided. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

If the operating environment for the system, storage, or associated cabling has changed in terms of temperature, humidity, vibration, unstable electrical power, or other physical factors, gracefully shut down and power off the system until the environment is restored to normal operations. If the operating environment has not changed, identify the source of the inconsistency by checking for previous errors and warnings. Also check for hardware statistics from Fibre Channel, SCSI, disk drives, other communications mechanisms, and previous administrative activities. Another approach is to consult with NetApp technical support (it is recommended that you do so before attempting this step), and then run WAFL® Iron (wafliron) on the aggregate to identify and optionally repair the inconsistency. If you need further assistance, contact NetApp technical support.

## Syslog Message

Call home for AGGREGATE INCONSISTENT AND TAKEN OFFLINE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.aggr.restricted

### Severity

ERROR

### Description

This message occurs when a medium error is encountered during a reconstruction. Medium errors during reconstruction automatically invoke waflliron if your system's 'raid.reconstruct.waflliron.enable' option is set to "on". If the option is set to "off", the aggregate is moved offline. In this case, the option was set to "off" and the aggregate was moved offline. The failed disk drive's reconstruction continues and the aggregate stays offline until you repair it. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run waflliron manually to bring this aggregate online. It is recommended that you contact NetApp technical support before attempting this. If you need further assistance, contact NetApp technical support.

## Syslog Message

Call home for MEDIUM ERROR DURING RECONSTRUCTION (aggregate restricted)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.aggr.waflliron

### Severity

ERROR

### Description

This message occurs when a medium error is encountered during a reconstruction. Medium errors during reconstruction automatically invoke waflliron if your system's 'raid.reconstruct.waflliron.enable' option is set to "on". If the option is set to "off", the aggregate is moved offline. In this case, the option was set to "on" and the aggregate remains online. The failed disk drive's reconstruction continues and waflliron begins repairing the aggregate. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The corrective process of running waflliron has already begun. Wait until you see a message indicating that the waflliron process is complete and the inconsistency status of the aggregate is cleared. If this doesn't happen, or if you see subsequent error messages, contact NetApp technical support. If you need assistance in bringing the aggregate online, contact NetApp technical support.

## Syslog Message

Call home for MEDIUM ERROR DURING RECONSTRUCTION (wafliron started)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.andu events

## callhome.andu.completed

### Severity

NOTICE

### Description

This message occurs when the automated update of the cluster has been completed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.andu.pausederr

### Severity

ALERT

### Description

This message occurs when the automated update of the cluster has been paused due to an error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The automated update of the cluster has been paused due to an error. Determine the cause of the error by using the "cluster image show-update-progress" command. Correct the cause of the error, and then issue the "cluster image resume-update" command to continue the automated update of the cluster.

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multinode AutoSupport.

# callhome.appdm events

## callhome.appdm.snaprest.fail

### Severity

ALERT

### Description

This message occurs when the Application Data Management feature encounters an error while attempting to restore an application from a Snapshot copy. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "application snapshot show -vserver <vserver> -application <app> -protection-group <pg>" command to determine if the snapshot exists. If it does, retry the operation. Contact NetApp technical support for assistance with reviewing the event log to determine the cause of the application snapshot restore failure.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**vserver\_name** (STRING): Vserver name related to this event.

**application\_name** (STRING): Application name related to this event.

**protection\_group\_name** (STRING): Protection group name related to this event.

**snapshot\_name** (STRING): Snapshot name related to this event.

# callhome.arw events

## callhome.arw.activity.seen

### Severity

ALERT

### Description

This message occurs when ransomware activity is detected. To protect the data, a Snapshot copy has been created, which can be used to restore the original data. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Refer to the anti-ransomware documentation to take remedial measures for ransomware activity. If you need assistance, contact NetApp technical support.

### Syslog Message

Call-home message for %s



## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**volName** (STRING): Name of the volume.

**volUuid** (STRING): UUID of the volume.

**vserverName** (STRING): Name of the Vserver.

**vserverUuid** (STRING): UUID of the Vserver.

# callhome.bad events

## callhome.bad.ram

### Severity

ERROR

### Description

This message occurs when correctable ECC errors are masked off, possibly due to bad or poorly seated memory or a faulty motherboard. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If anyone recently installed memory, try removing and reseating it. If problems persist, replace the memory module. If problems still persist, replace the motherboard. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for POSSIBLE BAD RAM

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.battery events

## callhome.battery.failure

### Severity

EMERGENCY

### Description

This message occurs when the battery is either incapable of maintaining data during a power-outage, or is experiencing a failure that poses a physical threat. To prevent data loss and/or damage, the system will shut down momentarily. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Correct any environmental issues, such as a high ambient temperature. If the condition persists, replace the NVRAM card/battery. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s CRITICAL.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.battery.low

### Severity

ALERT

### Description

This message occurs when the NVRAM battery capacity is critically low. If your system is configured to do so, it generates and transmits an AutoSupport® (or 'call home') message to NetApp® technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "system node environment sensors show" command to view the battery's current status, capacity, and charging state. If the battery was recently replaced or the system was nonoperational for an extended period of time, monitor the battery to verify that it is charging properly. If the battery run time continues to decrease below critical levels, the storage system will automatically shut down. Contact NetApp technical support for assistance.

## Syslog Message

Call home for BATTERY\_LOW.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## callhome.battery.notice

### Severity

NOTICE

### Description

This message occurs when the battery is operating near the limit of its supported environment. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Correct any environmental issues, such as a high ambient temperature. If the condition persists, replace the NVRAM card/battery in the next three months. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s NOTICE.

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.battery.overchg****Severity**

ALERT

**Description**

This message occurs when the NVRAM card lithium battery is overcharged and is at risk of failure. The system is currently serving data but will shut down shortly. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Replace the NVRAM card/battery and start the system back up. The system will not boot with the battery in this state. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for BATTERY OVERCHARGED

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.battery.overtemp****Severity**

ALERT

**Description**

This message occurs when the sensor for the NVRAM card lithium battery over the intended temperature limit. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

If the system is excessively warm, allow it to cool gradually. If the battery temperature is still too high, replace the battery pack. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for BATTERY OVERTEMP

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.battery.warning****Severity**

ALERT

## Description

This message occurs when the battery is operating outside the limit of its supported environment. Under these conditions, the battery might not be capable of maintaining data during a power-loss event. To prevent data loss, the system will automatically shut down in 24 hours. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Correct any environmental issues, such as a high ambient temperature. If the condition persists, replace the NVRAM card/battery. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s WARNING.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.bcom events

## callhome.bcom.vsinit.failed

### Severity

ALERT

### Description

This message occurs when multiple attempts to initialize the SAN Kernel Agent for a Vserver have failed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "cluster kernel-service show" (advanced) command to verify the cluster kernel quorum service is "in-quorum" on the source node. Contact NetApp technical support for assistance determining the root cause of the failed Vserver initialization.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**source** (STRING): The name of the node that failed to initialize.  
**vserver** (STRING): The name of the Vserver that failed to initialize.  
**vserver\_uuid** (STRING): The UUID of the Vserver that failed to initialize.  
**error** (STRING): The error of the last failed initialization.

# callhome.bdfu events

## callhome.bdfu.suspended

### Severity

ALERT

### Description

This message occurs when drive unlock fails for firmware download on a FIPS-certified NetApp Storage Encryption SAS drive.

### Corrective Action

Contact NetApp technical support for assistance with restoring valid authentication keys on the disk.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.bootfs events

### callhome.bootfs.chkdisk

### Severity

ALERT

### Description

This message occurs when a consistency check of the boot media (compact flash or other type) has shown a problem and has not been corrected automatically. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The boot media should be reformatted and/or replaced. After corrective steps have been taken, the "bootfs chkdisk" command should be used to confirm the recovery has been successful. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for BOOTFS CHKDSK

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.bridge events

### callhome.bridge.reboot.unexp

### Severity

ALERT

## Description

This message occurs when a bridge reboots unexpectedly. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Analyze the bridge event log to determine the cause of the reboot. If you need further assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**bridgeName** (STRING): Name of the bridge.

**bridgeVendor** (STRING): Vendor name of the bridge.

**bridgeModel** (STRING): Model name of the bridge.

**bridgeRevision** (STRING): Firmware revision of the bridge.

**bridgeSerialNum** (STRING): Serial number of the bridge.

**rebootReason** (STRING): Reason for the bridge reboot.

# callhome.c events

## callhome.c.fan

### Severity

ERROR

### Description

This message occurs when there is an unexpected decrease in a chassis fan's spin rate, possibly due to a slight decrease in revolutions or to a failed chassis fan. A failed chassis fan can lead to overheating and can cause damage to your equipment. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check the chassis fan to determine whether it is spinning. Replace the chassis fan if it failed or is not operating correctly. You might need to replace the motherboard to replace the chassis fan. If a slight decrease in revolutions seems to have caused this event, and if this has not happened often, you do not need to replace the chassis fan. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.c.fan.fru.degraded

### Severity

ERROR

### Description

This message occurs when a chassis fan might be operating in a degraded state. This condition can lead to overheating and can cause damage to your equipment. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check the chassis fan to determine whether it is spinning. Replace the chassis fan if it failed or is not operating correctly. You might need to replace the motherboard to replace the chassis fan. A slight decrease in revolutions might have caused this event. If this seems to be the case and has not happened often, you do not need to replace the chassis fan. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.c.fan.fru.fault

### Severity

ERROR

### Description

This message occurs when a chassis fan might have failed. This failure can lead to overheating and can cause damage to your equipment. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check the chassis fan to determine whether it is spinning. Replace the chassis fan if it failed or is not operating correctly. You might need to replace the motherboard to replace the chassis fan. A slight decrease in revolutions might have caused this event. If this seems to be the case and has not happened often, you do not need to replace the chassis fan. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.c.fan.fru.rm

**Severity**

ERROR

**Description**

This message occurs when a chassis fan is removed. You have two minutes to install a replacement chassis fan or the system shuts down automatically to avoid overheating. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Install a replacement chassis fan within two minutes of removing the original to complete the hot-swap. Otherwise, install the replacement chassis fan when you are ready. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.c.fan.fru.shut****Severity**

ERROR

**Description**

This message occurs when multiple chassis fans have failed. This failure can lead to overheating and can cause damage to your equipment. It is sent just before the system shutdown. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Check the chassis fans to determine whether they are spinning. Replace the chassis fans if any have failed or they are not operating correctly. You might need to replace the motherboard to replace the chassis fans. A slight decrease in revolutions might have caused this event. If this seems to be the case and has not happened often, you do not need to replace the chassis fan. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.carrier events****callhome.carrier.fault**



**Severity**

ERROR

**Description**

This message occurs when all the disks in the carrier, that contains more than one disk in it, have failed. The carrier should be replaced. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate the carrier with the failed disks by looking for the carrier whose amber LED is on and the green LED does not indicate any disk activity. Replace the carrier with the failed disks. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ch events

### callhome.ch.ps.degraded

**Severity**

ERROR

**Description**

This message occurs when a degraded chassis power supply unit (PSU) state exists. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

You might need to replace a chassis PSU. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.ch.ps.fan.bad.xmin

**Severity**

EMERGENCY

**Description**

This message occurs when a power supply fan has been detected as no longer working and the overheating is expected shortly. If your system is configured to do so, it generates and transmits an

AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Replace the power supply unit or the power supply unit's fan. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ch.ps.fault****Severity**

ERROR

**Description**

This message occurs when one chassis power supply unit (PSU) fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

You might need to replace a chassis PSU. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ch.ps.wrong.in****Severity**

ERROR

**Description**

This message occurs when a power supply is not connected to the correct type of power source (voltage, amperage, type, etc. difference). If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Connect the power supply to the proper power source or replace the power supply with the appropriate one for the available power source. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.chassis.events

## callhome.chassis.hitemp

### Severity

ERROR

### Description

This message occurs when the chassis temperature is too warm and might cause an environmental problem, or when there is faulty hardware in the system, such as a fan failure or a malfunctioning temperature sensor. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for CHASSIS OVER TEMPERATURE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.chassis.overtemp

### Severity

EMERGENCY

### Description

This message occurs when the chassis temperature is too hot. It is sent just before system shutdown and indicates a potential environmental problem or faulty hardware in the system, such as a fan failure or a malfunctioning temperature sensor. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If the environment is too hot, shut down and power off your system until the environmental conditions are restored to the normal operating temperature range. If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If

you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CHASSIS OVER TEMPERATURE SHUTDOWN

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

### **callhome.chassis.power**

### **Severity**

ERROR

### **Description**

This message occurs when power is in a degraded state because of bad power supply units (PSUs), bad wall power, or bad components on the motherboard. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

If spare PSUs are available, try exchanging them to see whether that alleviates the problem. If not, you might need to replace motherboard components or the motherboard. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

### **callhome.chassis.power.down**

### **Severity**

ERROR

### **Description**

This message occurs when voltage and other power sensors for the system power show critical values. A power supply unit (PSU) might have failed or the sensors might be malfunctioning. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Replace the faulty hardware. Most likely, you need to replace a PSU. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.chassis.ps.degraded****Severity**

ERROR

**Description**

This message occurs when a degraded chassis power supply unit (PSU) state exists. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

You might need to replace a chassis PSU. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.chassis.ps.fault****Severity**

ERROR

**Description**

This message occurs when one chassis power supply unit (PSU) fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

You might need to replace a chassis PSU. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.chassis.ps.off****Severity**

ERROR

**Description**

This message occurs when the system detects that a chassis power supply unit (PSU) is turned off. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to

NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Turn on the chassis PSU to restore normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.chassis.ps.ok**

### **Severity**

NOTICE

### **Description**

This message occurs when the system detects that a chassis power supply unit (PSU) was restored to normal operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

(None).

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.chassis.ps.replace**

### **Severity**

ERROR

### **Description**

This message occurs when a chassis power supply unit (PSU) fails. The PSU's DC power output was abnormal although the AC input power was within acceptable limits. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Replace the affected chassis PSU. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.chassis.ps.rm****Severity**

ERROR

**Description**

This message occurs when the system detects that a chassis power supply unit (PSU) was removed. You have two minutes to install a replacement chassis PSU or the system shuts down automatically to avoid overheating. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Install a replacement chassis PSU within two minutes of removing the original to complete the hot-swap. Otherwise, install the replacement chassis PSU when you are ready. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.chassis.undertemp****Severity**

EMERGENCY

**Description**

This message occurs when the chassis temperature is too cold. It is sent just before system shutdown and indicates a potential environmental problem or faulty hardware in the system, such as a fan failure or a malfunctioning temperature sensor. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If the environment is too cold, shut down and power off your system until the environmental conditions are restored to the normal operating temperature range. If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for CHASSIS UNDER TEMPERATURE SHUTDOWN

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.checksum events

## callhome.checksum

### Severity

ALERT

### Description

This message occurs when the system detects a RAID level checksum error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for CHECKSUM ERROR

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.checksum.multiple

### Severity

ALERT

### Description

This message occurs when the system detects a RAID checksum error involving multiple disk drives. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for CHECKSUM ERROR (multiple disks)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.clam events

## callhome.clam.node.ooq



**Severity**

EMERGENCY

**Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveness, Availability Monitor (CLAM) determines that one of the nodes is out of "cluster quorum".

**Corrective Action**

One or more of the following actions may be needed: - Check if the node that is out of quorum is administratively down; if it is, enable it with: "cluster modify -node node\_name -eligibility true"; - Check if the node that is out of quorum has panicked or hung; if it is, troubleshoot and resolve that problem. - Check the connectivity of the node that is out of quorum to the cluster interconnect. Fix any issues, if needed. - Check if there is a version incompatibility between the node that is out of quorum and the rest of the nodes in the cluster; if there is, upgrade nodes as needed so that they are compatible.

**Syslog Message**

Call home for %s.

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.client events

### callhome.client.app.alert

**Severity**

ALERT

**Description**

This message occurs when a remote client accesses the node using Manage ONTAP® and triggers an event using the ems-autosupport-log API call with "auto-support" field set to "true". If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Identify and resolve the issue on the remote client that is using the Manage ONTAP API to trigger this event. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.client.app.emerg

**Severity**

EMERGENCY

## Description

This message occurs when a remote client accesses the node using Manage ONTAP® and triggers an event using the `ems-autosupport-log` API call with "auto-support" field set to "true". If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Identify and resolve the issue on the remote client that is using the Manage ONTAP API to trigger this event. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.client.app.err

### Severity

ERROR

## Description

This message occurs when a remote client accesses the node using Manage ONTAP® and triggers an event using the `ems-autosupport-log` API call with "auto-support" field set to "true". If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Identify and resolve the issue on the remote client that is using the Manage ONTAP API to trigger this event. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.client.app.info

### Severity

NOTICE

## Description

This message occurs when a remote client accesses the node using Manage ONTAP® and triggers an event using the `ems-autosupport-log` API call with "auto-support" field set to "true". If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Identify and resolve the issue on the remote client that is using the Manage ONTAP API to trigger this event. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.client.app.notice

### Severity

NOTICE

### Description

This message occurs when a remote client accesses the node using Manage ONTAP® and triggers an event using the `ems-autosupport-log` API call with "auto-support" field set to "true". If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Identify and resolve the issue on the remote client that is using the Manage ONTAP API to trigger this event. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.clus events

### callhome.clus.net.degraded

### Severity

ALERT

### Description

This message occurs when the cluster network has experienced an issue that could potentially affect connectivity, reliability, redundancy, or performance, as observed from the local node. A degraded cluster network might cause a cluster-wide outage. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

See the EMS event log for additional information related to the issue. If the `event_type` is "Frequent Link Flapping", check the network cables that are connected to the specified port for any loose connections. If the `event_type` is "Large MTU Packet Loss", verify cluster switch MTU settings and ensure that all cluster

network ports are configured with the same MTU. If the event\_type is "Total Packet Loss", verify all cluster network connectivity and configuration and ensure that all cluster network ports are reachable from one another. If the event\_type is "Cluster LIF Not Assigned to Any Port" and the cluster ports are link "down", check the network cables that are connected to the specified port for any loose connections. Ensure that the cluster ports are administratively configured link "up". If the cluster ports are link "up", attempt to manually migrate the cluster LIF to another cluster port. If the event\_type is "Cluster LIF Administratively Down", use the "network interface modify" command to change the status-admin state of the cluster LIF to "up". If the event\_type is "Insufficient L2 Reachability", verify that the network topology is correct and the ports are correctly assigned to the right broadcast domain. If the problem persists after correcting any issues, call NetApp technical support.

### Syslog Message

Call home for %s: %s - %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**event\_type** (STRING): The type of the event that signaled that the cluster network might be operating in a degraded state from the perspective of the local node.

**event\_details** (STRING): Details of the event that signaled that the cluster network might be operating in a degraded state from the perspective of the local node.

## callhome.clus.vol.cre.fail

### Severity

ERROR

### Description

This message occurs when the active volume required to store the system data needed for cluster-wide storage cannot be created. This is typically due to not finding an aggregate with the specified size or mirrored attribute to create the volume. The specified feature cannot function without this volume. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Study the error message and take steps to rectify the problem. For example, if the error message indicates that the volume create failed because the aggregate is offline, use the 'aggr online -aggregate' command to bring it back online. If the error message does not indicate an obvious problem, run the 'aggregate show' command to find the state of the aggregate and 'metrocluster check config-replication show-aggregate-eligibility' command to find the eligibility of aggregates for configuration replication. If the aggregate is full, try to create space on it. Otherwise, if possible, create an aggregate with the specified mirror attribute and size by using the 'aggregate create -aggregate name -diskcount diskcount -disksize size -mirror mirrorattribute' command. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s. Could not create volume after %d attempts for feature %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**AttemptsMade** (INT): Count of failed attempts to create the volume.

**Feature** (STRING): Feature using the cluster-wide storage.

**ErrorMessage** (STRING): Error message on create failure at time of event generation

**Size** (STRING): Size of the volume to be created.

**Mirror** (STRING): Mirrored attribute of the aggregate.

## callhome.clus.vol.del.fail

### Severity

ERROR

### Description

This message occurs when the volumes that were used for cluster-wide storage cannot be deleted. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Study the error message and take steps to rectify the problem. For example, if the error message indicates that the volume delete failed because the aggregate is offline, use the 'aggr online -aggregate aggregate\_name' command to bring it back online. If the error message does not indicate an obvious problem, contact NetApp technical support

### Syslog Message

Call home for %s. Could not delete volumes "%s" and "%s" for feature %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**ActiveVolume** (STRING): Active volume name.

**StandbyVolume** (STRING): Standby volume name.

**Feature** (STRING): Feature using the cluster-wide storage.

## callhome.config events

### callhome.config

### Severity

ALERT

### Description

This message occurs when the system detects a disk shelf configuration error for one or more disk shelves. The system continues to report this error until you resolve the problem. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the command "storage errors show" to display outstanding storage errors in the system. Use the "event log show" command to identify earlier EMS messages related to shelf configuration errors. Correct the cause of the configuration error. Common causes of configuration errors are as follows: - A previously attached disk shelf was removed - A loss of disk shelf enclosure services occurred - A disk shelf fan module failed - An invalid disk shelf drive configuration was created - An operating environment temperature was 50C (122F) degrees or higher - A disk shelf environment module failure occurred If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for CONFIGURATION\_ERROR

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.coredump events

## callhome.coredump.save.done

### Severity

NOTICE

### Description

This message occurs when all the core segments belonging to a coredump have been saved. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "system node coredump show" command to obtain coredump file details.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**file** (STRING): Primary core segment file that was created.

**node** (STRING): Name of the node whose core was saved.

**size** (LONGINT): Size of primary core segment in bytes.

**chksum** (STRING): Checksum of primary core segment file's data.

**ancillary\_segment\_directory** (STRING): Directory which holds ancillary core segments.

**ancillary\_segments** (STRING): List of ancillary segments.

# callhome.corefile events

## callhome.corefile.upload

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport for corefile upload. This Autosupport collects and sends only the corefile specified as part of the payload. This autosupport is sent only when the transport is set to http or https. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

## Syslog Message

corefile upload callhome for %s

### Parameters

**subject** (STRING): AutoSupport message subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport message should be generated and delivered even if the AutoSupport configuration is disabled. If "0" then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this coredump AutoSupport message.

**core\_filename** (STRING): Coredump filename to be included in the coredump AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport message via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

## callhome.cpu events

### callhome.cpu.fan

#### Severity

ERROR

#### Description

This message occurs when a CPU fan failed or the fan is spinning too slowly or too fast. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Analyze whether you need to replace the CPU fan. CPU fan replacement might require that you replace the motherboard as well. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for %s

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.cpu.fan.fault

#### Severity

ERROR

#### Description

This message occurs when a CPU fan failed or the fan is spinning too slowly or too fast. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Analyze whether you need to replace the CPU fan. CPU fan replacement might require that you replace the motherboard as well. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.css events

### callhome.css.toc.repair

#### Severity

ALERT

#### Description

This message occurs when an inconsistency is discovered in a metadata record of the Cloud Storage Service (CSS). The inconsistency is automatically repaired immediately upon detection. But it can leave behind lost objects that are no longer pointed to by the namespace. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

The (privilege: diag) "object-store-server lost-object-recovery start" command can be used to add lost objects back into the namespace. The progress of this operation can be monitored with the "object-store-server lost-object-recovery show" command. If you need further assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ctran events

### callhome.ctran.gcror.timeout

#### Severity

ALERT

#### Description

This message occurs when the Kernel Cluster Services (KCS) / Cluster Transaction Manager (CTRAN) detects that one or more newly selected Group Coordinators (GCs) are taking too long to successfully complete the GC Ready on Reselect procedure. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.



### **Corrective Action**

Call NetApp technical support immediately for further assistance.

### **Syslog Message**

Call home for GROUP COORDINATOR ROR TIMEOUT.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## **callhome.ctran.jpc.splitbrain**

### **Severity**

ALERT

### **Description**

This message occurs when the Kernel Cluster Services (KCS) / Cluster Transaction Manager (CTRAN) detects that there is a disagreement in the cluster - a split brain - about which node has the Join Proposal Coordinator (JPC). If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Call NetApp technical support immediately for further assistance.

### **Syslog Message**

Call home for CTRAN JPC SPLIT BRAIN DETECTED.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## **callhome.dafs events**

### **callhome.dafs.nic.stop**

### **Severity**

ALERT

### **Description**

This message occurs when the DAFS (Direct Access File System) NIC (Network Interface Controller) has stopped while DAFS was running. This is caused by a failure of the NIC processor firmware during heavy and complex workloads. DAFS can usually recover and continue. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Reboot the system to recover the NIC. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DAFS VI NIC HALTED - DAFS SERVICE MAY BE INACCESSIBLE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.dbs events

## callhome.dbs.assign.error

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) assignments reports a validation error. An assignment validation error may indicate that data is inaccessible or data loss has occurred. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for assistance.

## Syslog Message

Call home for %s, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**evtMessage** (STRING): Description of the data event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It will be "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It will be 0 if there is no associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It will be 0 if there is no associated drive.

**driveUuid** (STRING): Drive UUID of the associated drive. It will be zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**evtDetails** (STRING): Specific details of the DBS data event. It may be a string, in JSON format, or empty.

## callhome.dbs.block.repair

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) attempts a new block repair or fails to repair block in maximum number of repair attempts allowed. See event details for more information. The DBS is responsible for managing the data that backs the FlexVols®. Multiple DBS events that indicate inconsistent

data are mapped to this one event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for help immediately.

### Syslog Message

Call home for %s, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**evtMessage** (STRING): Description of the data event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It will be "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It will be 0 if there is no associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It will be 0 if there is no associated drive.

**driveUuid** (STRING): Drive UUID of the associated drive. It will be zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**evtDetails** (STRING): Specific details of the DBS data event. It may be a string, in JSON format, or empty.

## callhome.dbs.cannot.sync.blk

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects the cluster Block data is in a degraded state, and the auto-heal process to restore full Block data redundancy cannot proceed. Block resources need to be added or restored to the cluster. Too many nodes, Block services, FireStorm services or aggregates may be offline; or possibly the cluster Block services are full. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Evaluate the storage capacity of the cluster. If some capacity is currently unavailable due to drives or nodes being offline, attempt the repairs needed to bring them online. Free storage used by FlexVols by deleting unneeded volumes.

### Syslog Message

The DBS cluster Block data cannot be made fully redundant.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.clcap.full

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects the cluster (or node in the cluster) has completely consumed its capacity. The DBS is responsible for managing the data that backs the FlexVols®. The cluster data is now read-only. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Evaluate the available capacity of this cluster relative to its usage. If some capacity is currently unavailable due to drives or nodes being offline, attempt the repairs needed to bring them online. Free storage used by FlexVols by deleting unneeded volumes. If the cluster's available capacity appears inconsistent with the calculated expected capacity, then contact NetApp technical support.

### Syslog Message

The DBS cluster capacity is 100%% full.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.clcap.near.full

### Severity

ALERT

## Description

This message occurs when the Distributed Block Store (DBS) detects the cluster (or node in the cluster) has consumed nearly all of its capacity. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Evaluate the available capacity of this cluster relative to its usage. If some capacity is currently unavailable due to drives or nodes being offline, attempt the repairs needed to bring them online. Free storage used by FlexVols by deleting unneeded volumes.

## Syslog Message

The DBS cluster capacity is nearly full.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.clst.ekm.server

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects one or more of the following: 1. An External Key Management server cannot be reached by any node and is considered down. 2. A fatal communication error has occurred with an External Key Management server. Seek corrective-action immediately.

### Corrective Action

Verify that the External Key Management server is running and can be reached by the management network. Verify that the correct root CA certificate and client certificate are configured for the Key Server (use openssl s\_client to assist with troubleshooting). Check the logs on the External Key Management server for additional troubleshooting help. Contact NetApp technical support for assistance.

### Syslog Message

Callhome for %s. Fault type %s, fault ID %u, status %u, nodes %s. %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.data.unavail

### Severity

EMERGENCY

### Description

This message occurs when some portion of the Distributed Block Store (DBS) data cannot be accessed. The DBS is responsible for managing the data that backs the FlexVols®. Multiple DBS events that indicate data is unavailable are mapped to this one event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for help immediately.

### Syslog Message

Call home for %s, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**evtMessage** (STRING): Description of the data event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It will be "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It will be 0 if there is no associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It will be 0 if there is no associated drive.

**driveUuid** (STRING): Drive UUID of the associated drive. It will be zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**evtDetails** (STRING): Specific details of the DBS data event. It may be a string, in JSON format, or empty.

## callhome.dbs.dup.master

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects an internal database inconsistency. Master services have multiple entries for cluster master candidacy. Manual intervention is required.

## Corrective Action

Contact NetApp technical support immediately.

## Syslog Message

Callhome for %s. DBS has detected a critical issue. Cluster fault type %s, fault ID %u, status %u, nodes %s. %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Details of the duplicate master candidates.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ekm.cert.emer

## Severity

EMERGENCY

## Description

This message occurs when the Distributed Block Store (DBS) detects that an External Key Management server configuration contains a certificate that is invalid, is expired, or will expire in less than 3 days. Seek corrective-action immediately.

## Corrective Action

Renew or correct each listed certificate and update the associated Key Server configuration. Contact NetApp technical support for assistance.

## Syslog Message

Callhome for %s. Fault type %s, fault ID %u, status %u, nodes %s. %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ensemble.emer

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) loses the database connection to over half of the ensemble nodes (of 3 or 5 total). DB quorum has been lost and the DB is down. The cluster cannot operate without access to the DB.

### Corrective Action

Bring offline nodes back online and/or restore node network connectivity. If that does not resolve the fault, contact NetApp technical support.

### Syslog Message

Callhome for %s. DBS has lost network connectivity or power on ensemble node. Cluster fault type %s, fault ID %u, status %u, nodes %s. %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the degraded ensemble.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ffmAggrFlagSet

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects one or more aggrs have a don't attach, don't mount or don't detach flag set.

### Corrective Action

Resolve the issue that required the flag to be set, then remove the flag using the SetAggrFlags API command.

### Syslog Message

Callhome for %s. Fault type %s, fault ID %u, status %u, nodes %s. %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.



**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ffmFlexvolsOffl

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects one or more FlexVols are offline.

### Corrective Action

Solve any issue contributing to the FlexVols being taken offline.

### Syslog Message

Callhome for %s. Fault type %s, fault ID %u, status %u, nodes %s. %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.generic.cfault

### Severity

EMERGENCY

### Description

This message occurs when a Distributed Block Store (DBS) emergency cluster fault is reported as a generic fault. The DBS is responsible for managing the data that backs the FlexVols®. Multiple faults are mapped to this one event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for interpretation of this event.

## Syslog Message

Call home for %s, fault %s (%u), sev %u, service %u, node %u/%s, drives %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**cfCodeName** (STRING): Name of the original DBS cluster fault code.  
**cfCode** (INT): DBS cluster fault code number.  
**cfSeverity** (INT): Severity of the original DBS cluster fault. This is different than the EMS severity.  
**serviceID** (INT): Service ID that identifies the associated cluster service. It will be "0" if there is no associated service.  
**nodeID** (INT): Node ID number of the associated node. It will be "0" if there is no associated node.  
**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.  
**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.  
**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Description of the cluster fault, including context details.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.inv.block.data

### Severity

EMERGENCY

### Description

This message occurs when a block stored in the Distributed Block Store (DBS) has inconsistent data. The cluster is checking the compression, block ID and checksum of the block against expected values. The DBS is responsible for managing the data that backs the FlexVols®. Multiple DBS events that indicate inconsistent data are mapped to this one event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for help immediately.

## Syslog Message

Call home for %s, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**evtMessage** (STRING): Description of the data event, including context details.  
**evtType** (STRING): Type of the original DBS event.  
**serviceID** (INT): Service ID that identifies the associated DBS service. It will be "0" if there is no associated service.  
**nodeID** (INT): Node ID number of the associated node. It will be 0 if there is no associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.  
**driveID** (INT): Drive ID number of the associated drive. It will be 0 if there is no associated drive.  
**driveUuid** (STRING): Drive UUID of the associated drive. It will be zeroes if there is no associated drive.  
**evtID** (INT): DBS event ID number.  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**evtDetails** (STRING): Specific details of the DBS data event. It may be a string, in JSON format, or empty.

## callhome.dbs.mem.thresh.emer

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects container memory 100 percent consumed.

### Corrective Action

Add additional nodes, or purge deleted volumes and run GC

### Syslog Message

Callhome for %s. DBS has detected free memory is completely consumed on a node's container. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u, nodes %s. %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**nodeID** (INT): Node ID number of the associated node.  
**nodeUUID** (STRING): Node UUID string of the associated node.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Details of the memory usage threshold.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.multSliceSvcDn

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects multiple unhealthy slice services. This may cause some data to be unavailable.

### Corrective Action

If multiple slice services are unhealthy, repair offline nodes and drives to bring slice services back online. Contact NetApp technical support.

## Syslog Message

Call home for %s. Fault type %s, fault id %u, status %u, nodes %s. %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Description of the cluster fault, including context details.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.node.ekm.server

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects that a node is unable to reach an External Key Management server that is not considered down (other nodes are able to reach it). Seek corrective-action immediately.

### Corrective Action

Perform troubleshooting at the network or node-specific level to determine why the listed node is unable to reach the External Key Management server. Contact NetApp technical support for assistance.

## Syslog Message

Callhome for %s. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u, nodes %s. %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**nodeID** (INT): Node ID number of the associated node.  
**nodeUUID** (STRING): Node UUID string of the associated node.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Description of the cluster fault, including context details.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.sec.cache.full

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects that Slice volume secondary write cache is full. This is caused when Slice service write requests to the Block service (and by extension the FireStorm service) are not receiving replies. Client write performance may be reduced while this condition exists. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Verify that all nodes are online. Verify that all Block and FireStorm services, and associated aggregates, are online. Attempt repairs needed to bring nodes and aggregates online. If the reason for this condition cannot be found, contact NetApp technical support.

### Syslog Message

Call home for %s, service %u, node %u/%s, drives %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**serviceID** (INT): Service ID that identifies the associated cluster service.  
**nodeID** (INT): Node ID number of the associated node.  
**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.  
**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.  
**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Description of the cluster fault, including context details.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ssl.ndcert.exc

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects that a node has a missing or corrupted certificate.

### Corrective Action

Contact NetApp technical support for assistance.

## Syslog Message

Callhome for %s. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u, nodes %s. %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**nodeID** (INT): Node ID number of the associated node.

**nodeUUID** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.ssl.ndcert.exp

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects that a node contains an SSL certificate that has expired or will expire imminently.

### Corrective Action

Contact NetApp technical support for assistance.

## Syslog Message

Callhome for %s. Cluster fault node %u, node UUID %s, fault type %s, fault ID %u, status %u, nodes %s. %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**nodeID** (INT): Node ID number of the associated node.

**nodeUUID** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.svcap.full

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects a service (Slice or Block) that has completely consumed its capacity. The DBS is responsible for managing the data that backs the FlexVols®. The data for this service is now read-only. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Evaluate the available capacity of this service relative to its usage. If some capacity is currently unavailable due to drives or nodes being offline, attempt the repairs needed to bring them online. Free storage used by FlexVols by deleting unneeded volumes. If the services's available capacity appears inconsistent with the calculated expected capacity, then contact NetApp technical support.

### Syslog Message

Call home for %s, service %u, node %u/%s, drives %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**serviceID** (INT): Service ID that identifies the associated cluster service.  
**nodeID** (INT): Node ID number of the associated node.  
**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.  
**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.  
**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.  
**cfDetails** (STRING): Description of the cluster fault, including context details.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.  
**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.svcap.near.full

### Severity

ALERT

### Description

This message occurs when the Distributed Block Store (DBS) detects a service (Slice or Block) that has consumed nearly all of its capacity. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Evaluate the available capacity of this service relative to its usage. If some capacity is currently unavailable due to drives or nodes being offline, attempt the repairs needed to bring them online. Free storage used by FlexVols by deleting unneeded volumes.

## Syslog Message

Call home for %s, service %u, node %u/%s, drives %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**serviceID** (INT): Service ID that identifies the associated cluster service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.

**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.dbs.volumes.offline

### Severity

EMERGENCY

### Description

This message occurs when the Distributed Block Store (DBS) detects that one or more volumes in the storage cluster are offline. For the volumes listed, both copies of volume slice data are unavailable. The volumes are degraded and the primary slice service becomes unresponsive, as detected by the Cluster Master. The VolumesDegraded fault will also be present. The DBS is responsible for managing the data that backs the FlexVols®. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check for network connectivity issues and hardware errors. There should be other faults if specific hardware components have failed. Attempt to resolve these issues. The fault will clear when the primary slice service is responsive. Contact Support if there are no associated hardware faults and this fault persists long enough to affect client applications.

### Syslog Message

Call home for %s. One or more DBS volumes offline.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.



**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**collectAsupNodes** (STRING): List of the node universal unique identifiers (UUIDs) from which AutoSupport messages should be collected.

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

**skip\_callhome** (INT): If 1, then the ASUP call home was not executed for this event. Call home is only done for newly reported faults (see cfStatus).

## callhome.device events

### callhome.device.qual

#### Severity

ALERT

#### Description

This message occurs when the system detected a disk drive device that did not receive qualification from NetApp. Unqualified disk drives can cause serious damage to the system, and the system shuts down within 72 hours to prevent this potential damage. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

To correct this situation, complete one of the following steps: 1. Download and install the latest Drive Qualification Package (DPQ) from NetApp technical support. This automatically qualifies any disk drives that were provided by NetApp. Search the support site knowledge base for the "DQP" keyword for further information. 2. Remove the unqualified disk drives from the system. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for DEVICE\_QUALIFICATION\_FAILED

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.disk events

### callhome.disk.carrier.evacuation.notpossible

#### Severity

ALERT

#### Description

This message occurs when a system does not contain enough spares to correct existing disk failures. This message is seen only on systems containing disk carriers capable of holding more than one disk. When this message occurs, manual intervention is required because ONTAP® cannot automatically correct the

existing issues with the system's disk configuration. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

To correct this issue on systems that contain carriers with multiple disks, try the following (in the specified order): (1) Assign any unassigned disks that are contained within a multidisk carrier. (2) Add multidisk carriers. (3) If a replacement carrier is available, do the following: a) Fail ALL disks within the carrier. Do not remove a multidisk carrier that contains non-failed disks; doing so results in system degradation and might result in system failure. b) Wait until the disk's amber Attention LEDs stop flashing and are lit continuously. Do not remove a carrier while the shelf OPS panel is alternating between the shelf ID and the Attention Code H1, or while the disks within the carrier are flashing their amber Attention LEDs. c) Replace the carrier. (4) Reassign partner spare disks in the HA configuration. (5) If the problem persists, call NetApp technical support.

### **Syslog Message**

Call home for DISK\_CARRIER\_EVACUATION\_NOTPOSSIBLE

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.disk.low.io**

### **Severity**

NOTICE

### **Description**

This message occurs when the storage health monitor detected low disk drive throughput. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

(None).

### **Syslog Message**

Call home for DISK\_LOW\_THRUPUT

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.disk.orphan**

### **Severity**

NOTICE

### **Description**

This message occurs when RAID subsystem detects an orphan or broken drive and is processing it. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for %s.

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.diskown events

### callhome.diskown.AssignExceeded

**Severity**

ERROR

**Description**

This message occurs when too many disks have been installed. An attempt is made to assign more disks than the maximum that can be supported on this node. This problem must be resolved before the system is rebooted. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Review the system capacity limits and the assigned disks on this system. If possible, unassign the unused array LUNs. If you need further assistance, contact NetApp technical support.

**Syslog Message**

Call home for ATTEMPTED TO ASSIGN MORE DISKS THAN MAXIMUM SUPPORTED.

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.disks events

### callhome.disks.missing

**Severity**

ERROR

**Description**

This message occurs when the system detected simultaneously that multiple disk drives were missing from the system. It generated a single "MULTIPLE DISKS MISSING" message rather than multiple "DISK MISSING" messages. The problem might be due to a Fibre Channel loop (connection) or a disk shelf or disk drive operating incorrectly. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for MULTIPLE DISKS MISSING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.dr events

## callhome.dr.apply.failed

### Severity

ERROR

### Description

This message occurs when all attempts to replicate the configuration fails. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To troubleshoot further, search the event management system (EMS) logs on the destination cluster for events starting with "vsdr". Perform any provided corrective actions. For further assistance, contact NetApp technical support.

## Syslog Message

Call home for DR CONFIG REPLICATION FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**DestinationVserver** (STRING): Destination Vserver of the DR relationship.

**FailedReason** (STRING): Failed reason.

# callhome.dsk events

## callhome.dsk.bypass

### Severity

ERROR

### Description

This message occurs when a disk drive failure occurred and the Embedded Switched Hub (ESH) removed the disk drive from the loop electronically—bypassing the disk drive on one or possibly both ports. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Depending on your system configuration, you might need to place the disk drive software into a failed state to match what the ESH module did. Then locate and replace the failed disk drive. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for DISK\_FAIL - Bypassed by ESH

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.dsk.config**

### **Severity**

ERROR

### **Description**

This message occurs when the system detects a configuration problem with a disk drive, Fibre Channel switch, or host bus adapter (HBA). A cabling issue might be causing the problem. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Review and repair the cabling and SAN configuration. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for DISK CONFIGURATION ERROR

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.dsk.disable**

### **Severity**

ERROR

### **Description**

This message occurs when a disk drive failure occurred and the I/O Module (IOM) removed the disk drive from the system by disabling the PHY on one or possibly both ports. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Depending on your system configuration, you might need to place the disk drive software into a failed state to match what the IOM did. Then locate and replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK\_FAIL - Disabled by IOM

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.fault

## Severity

ERROR

## Description

This message occurs when a disk drive fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the "storage disk show -broken" command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.fw.needed

## Severity

ERROR

## Description

This message occurs when the system detects one or more disk drives with out of date or "downrev" disk firmware. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Check and update all the disk drive firmware levels to the current the disk drive firmware level. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK\_FIRMWARE\_NEEDED\_UPDATE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.init

### Severity

ERROR

### Description

This message occurs when a disk drive failed to initialize. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "storage disk show -broken" command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for DISK INIT FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.io.degraded

### Severity

ERROR

### Description

This message occurs when the storage health monitor detects degraded disk I/O (taking too long) that indicates a problem with a disk drive. A disk drive might not be using the most recent disk drive firmware. If the disk drive is using the most recent firmware, it might be at a higher risk of failing. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Identify the disk drive that caused this event and verify that it is running the most recent firmware. If the firmware is not up to date, update it to ensure that problems do not continue. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for DISK\_IO\_DEGRADED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.label

### Severity

ALERT

## Description

This message occurs when a disk drive failed due to unexpected RAID metadata. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the “storage disk show -broken” command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK BAD LABEL

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.label.broken

## Severity

ERROR

## Description

This message occurs when a disk drive fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the “storage disk show -broken” command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK LABELED BROKEN

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.label.v

## Severity

ERROR

## Description

This message occurs when a disk drive failed due to a version number in its RAID metadata that is inconsistent with the operating system version installed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.



### **Corrective Action**

Use the “storage disk show -broken” command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for DISK BAD LABEL VERSION

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.dsk.missing**

### **Severity**

ERROR

### **Description**

This message occurs when the system cannot detect a disk drive in the Fibre Channel network (or loop) that it previously discovered. The error might be caused by a failed disk drive in the chassis still awaiting replacement or a broken connection to the disk drive. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Locate and replace the defective disk drive. If you already replaced it, check the Fibre Channel network for network connectivity from the appliance to the disk drive. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for DISK MISSING

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.dsk.no.spin**

### **Severity**

ERROR

### **Description**

This message occurs when a disk drive fails to start. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Logged message 'disk.init.failure.spinup' identifies the failed disk drive. Locate and replace the disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK NOT SPINNING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.noio

## Severity

ERROR

## Description

This message occurs when disk drive device commands to a disk drive are failing with link errors, adapter errors, or timeouts. With this type of problem, the disk drives's red or amber LED fault light may not be lit. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the "storage disk show -broken" command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK NOT RESPONDING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.predict

## Deprecated

Deprecated as of ONTAP 9.2. Last used in DOT 7.1.1.

## Severity

ALERT

## Description

This message occurs when the storage health monitor PFA (predictive failure analysis) determines that disk drive failure is forthcoming. A disk drive might not be using the most recent disk drive firmware. If the disk drive is using the most recent firmware, it might be at higher risk of failing. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Identify the disk drive that caused this event and check whether the firmware is up-to-date. If the firmware is not up-to-date, update it to ensure that problems do not continue. If the firmware is up-to-date, manually fail the faulty disk drive and remove it from the system to ensure the integrity of data. At the earliest opportunity, install a replacement disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK\_FAILURE\_PREDICTED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.redun.fault

## Severity

ERROR

## Description

This message occurs when redundant access to disk storage is lost. This message and associated Autosupport are issued at most once per day. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Locate disks which have lost redundancy by using the following command: "storage disk show -fields disk,primary-port,secondary-name,secondary-port,shelf,bay" Note all channels with single path connectivity. Check for failures using these commands: "system health node-connectivity shelf show" and "system health system-connectivity shelf show". If a module or cable failure is identified, replace the failed module or cable. Confirm redundancy has been restored with the following command: "storage disk show -fields disk,primary-port,secondary-name,secondary-port,shelf,bay" If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK REDUNDANCY FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.scrub

## Severity

ALERT

## Description

This message occurs when the weekly scheduled disk scrub detects checksum errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Review the results of the disk scrub. Your next steps depend on the result of your review, and might include replacing the disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK SCRUB CHECKSUM ERROR

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.dsk.scrub.done****Severity**

ERROR

**Description**

This message occurs when a disk scrub finishes bypassing media errors and places the aggregate in a normal state. This type of disk scrub is started automatically after wafiron when disk reconstruction is bypassing media errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Review the results of the disk reconstruction, wafiron, and disk scrub. Your next steps depend on the result of your review, and might include replacing the disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for DISK SCRUB COMPLETED (bypassing media errors)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.dsk.scrub.parity****Severity**

ALERT

**Description**

This message occurs when the scheduled disk scrub detects parity inconsistencies. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Review the results of the disk scrub. Your next steps depend on the result of your review, and might include replacing the disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for DISK SCRUB PARITY INCONSISTENCY

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.dsk.sense****Severity**

ERROR

## Description

This message occurs when storage health monitor detects a SCSI sense error that indicates a problem with a disk drive. A disk drive might not be using the most recent disk drive firmware. If the disk drive is using the most recent firmware, it might be at a higher risk of failing. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Identify the disk drive that caused this event and verify that it is running the most recent firmware. If the firmware is not up-to-date, update it to ensure that problems do not continue. If the firmware is up-to-date, manually fail the faulty disk drive and remove it from the system to ensure the integrity of data. At the earliest opportunity, install a replacement disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK\_RECOVERED\_ERRORS

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dsk.shrank

### Severity

ALERT

## Description

This message occurs when either an aggregate or hot spare disk drive fails because its raw capacity was downsized due to a firmware change. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the "storage disk show -broken" command to locate the failed drive. Confirm the location by looking for a drive whose red and amber LED fault lights are lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for DISK RAWSIZE SHRANK

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.dumptriage events

## callhome.dumptriage

### Severity

NOTICE

## Description

This message occurs when kernel dumptrriage files are collected for a coredump event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**core\_filename** (STRING): Coredump filename used to find the correct set of dumptrriage files.

# callhome.encl events

## callhome.encl.access

### Severity

ERROR

### Description

This message occurs when contact to one or more disk shelf enclosure services (environmental monitoring) is lost, possibly due to storage cabling, adapter, or configuration problems. The system uses the enclosure services to monitor the environment continuously to ensure that conditions remain at safe operating levels. If the system cannot contact the enclosure services, it cannot detect when operating conditions move into unsafe ranges and cannot minimize any damage that might result. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "storage errors show" command to display the current outstanding shelf error conditions. Use the "event log show" command to identify earlier EMS messages related to shelf connectivity errors. Restore connectivity to the disk shelf enclosure services. Check cabling for loose connections. Swap out adapters, if possible. Double-check the configuration. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for ENCLOSURE SERVICES ACCESS ERROR

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.ext events

## callhome.ext.cache.bad

**Severity**

ALERT

**Description**

This message occurs when a check of the external memory cache used for increasing system performance has failed and this cache has been disabled. System operations will continue with possibly reduced performance. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

The external memory cache should be checked and possibly repaired to restore its service to the system. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for EXTERNAL CACHE FAILURE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fans events

### callhome.fans.failed

**Severity**

EMERGENCY

**Description**

This message occurs when multiple fans inside the appliance case fail. This failure can lead to overheating and cause damage to your equipment. The system shuts down within 10 minutes unless you resolve the fan problems. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Replace the fans within 10 minutes. If that is not possible, perform a graceful shutdown of the system and power it off. Then replace the fans when you can, power on the system, and restore it to normal operation. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for MULTIPLE FAN FAILURE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fc events

### callhome.fc.al.link

## Severity

ALERT

## Description

This message occurs when a Fibre Channel arbitrated loop (FC-AL) has link failures. You need to resolve this issue promptly. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Check the FC-AL itself and the connections to the FC-AL for poor or broken connections and misconfigurations. Examine the event log for messages related to FC or disk I/O errors and take the corrective actions found there. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fc.al.poor

## Deprecated

Deprecated as of ONTAP 9.2. Last used in DOT 7.1.0.

## Severity

ALERT

## Description

This message occurs when the system detects that a Fibre Channel arbitrated loop (FC-AL) is unreliable or is too poor in quality for data exchange. You need to resolve this issue promptly. The link errors are recoverable in the sense that the system is up and running. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Check the FC-AL itself and the connections to the FC-AL for poor or broken connections and misconfigurations. Examine the event log for messages related to FC or disk I/O errors and take the corrective actions found there. Restore normal, error-free communication for the FC-AL as quickly as possible. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for FC-AL RECOVERABLE ERRORS

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fcp events



## callhome.fcp.adapter.failed

### Severity

ERROR

### Description

This message occurs when a Fibre Channel Protocol (FCP) HBA Host Bus Adapter (HBA) in this node has failed to read a critical setting from the adapter's internal configuration. The adapter should be replaced. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Locate the FCP adapter in the slot specified in this message. Replace the adapter. If the adapter name starts with 0, the motherboard may need to be replaced. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fcp.sram.dump

### Severity

ERROR

### Description

This message occurs when a Fibre Channel Protocol (FCP) HBA Host Bus Adapter (HBA) in this node has recorded to disk a memory image from its Static Random Access Memory (SRAM). The SRAM dump has been performed as the result of several possible reasons, including a command directive and HBA firmware restart. Note that performing a dump of a HBA can disrupt data access operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If this event happens as a result of a command directive such as "fcp dump", the dump operation is expected. If this event happens as a result of HBA reset, the collected dump files can be used by technical support for subsequent analysis. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for FCP SRAM DUMP.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**portInstance** (INT): Instance number of the target port where the dump occurs.

**dumpDirName** (STRING): Directory location where the FC dumps are stored.

**portName** (STRING): FC HBA's name.

# callhome.fdisk events

## callhome.fdisk.admin

### Severity

NOTICE

### Description

This message occurs when you fail a file system disk drive. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for FILESYSTEM DISK ADMIN FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fdisk.fault

### Severity

ERROR

### Description

This message occurs when a disk drive fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fdisk.label

### Severity

ERROR

### Description

This message occurs when a disk drive fails due to unexpected RAID metadata. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp

technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for FILESYSTEM DISK BAD LABEL

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.fdisk.label.v**

### **Severity**

ERROR

### **Description**

This message occurs when a disk drive fails due to a version number in its RAID metadata that is inconsistent with the operating system version installed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for FILESYSTEM DISK BAD LABEL VERSION

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.fdisk.missing**

### **Severity**

ERROR

### **Description**

This message occurs when the system detects that a disk drive is not in a loop or is missing. This might also occur when a failed disk drive is still in the chassis awaiting replacement. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Locate and replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for FILESYSTEM DISK MISSING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fdisk.noio

## Severity

ERROR

## Description

This message occurs when disk drive device commands to a file system disk drive are failing with link errors, adapter errors, or timeouts. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for FILESYSTEM DISK NOT RESPONDING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fdisk.shrank

## Severity

ERROR

## Description

This message occurs when an aggregate disk drive fails because its raw capacity was downsized due to a firmware change. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. After the automated e-mail is sent to you with the case# and instructions, the case is closed. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for FILESYSTEM DISK RAWSIZE SHRANK

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.flash events

## callhome.flash.cache.degraded

### Severity

ERROR

### Description

This message occurs when Flash Management Module (FMM) detects that a percentage of a Flash Cache is offline that is above a warning threshold. FMM monitors all flash devices in the system that attach with it. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

A Flash Cache has degraded. This is an indication or the cause of a performance degradation. Exact impact cannot be estimated. This Flash Cache needs to be repaired or replaced. Contact customer support for more details.

### Syslog Message

Call home for Flash Cache DEGRADED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.flash.cache.failed

### Severity

ALERT

### Description

This message occurs when Flash Management Module (FMM) detects that a Flash Cache has suffered a failure. Typically, this is the result of a hardware failure on the Flash Cache itself. FMM monitors all flash devices in the system that attach with it. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The Flash Cache can fail due to many reasons. If you see "iomem.dma.internal" in EMS log, reboot the system. If the state of the Flash Cache in system configuration shows 'Failed firmware', update the device (follow the corrective action of the "iomem.card.fail.fpga.rev" EMS message). For all other reasons, replace the Flash Cache. For further assistance, contact technical support.

### Syslog Message

Call home for Flash Cache FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.flash.cache.offlined

### Severity

ALERT

### Description

This message occurs when Flash Management Module (FMM) detects that a critical percentage of a Flash Cache is offline, beyond which the device cannot operate. FMM turned off this device. FMM monitors all flash devices in the system that attach with it. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

A Flash Cache was offlined. This is an indication or the cause of a performance degradation. Exact impact cannot be estimated. This Flash Cache needs to be repaired or replaced. Contact customer support for more details.

### Syslog Message

Call home for Flash Cache OFFLINED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.fs events

### callhome.fs.lun.resized

### Severity

ALERT

### Description

This message occurs when the system has detected that a filesystem logical disk array or LUN (Logical Unit Number) stored raw size no longer matches the reported physical size of the LUN. Resizing of a filesystem gateway LUN is not supported. Only gateway LUNs that are spare (not used in a filesystem) may be resized. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Restore the logical disk array or LUN raw size to the expected file system size that system expects. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for FILESYSTEM LUN RESIZED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.hainterconnect events

## callhome.hainterconnect.down

### Severity

ALERT

### Description

This message occurs when the high-availability (HA) interconnect is down. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Corrective action depends on the number and type of HA interconnect links supported by the platform, as well as the reason for which the interconnect is down. EMS event "ic.HAInterconnectDown" provides the reason information. a. If links are down: Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that the small form-factor pluggables (SFPs), if applicable, are seated properly on both controllers. For internally connected links, disable and reenabte the links one at a time by using the "ic link off" and "ic link on" commands. b. If links are disabled: Enable the links by using the "ic link on" command. c. If a peer is not connected: Disable and reenabte the links one at a time by using the "ic link off" and "ic link on" commands. Contact NetApp technical support if the problem persists.

### Syslog Message

Call home for %s due to %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**reason** (STRING): Reason for interconnect down.

# callhome.hba events

## callhome.hba.failed

### Severity

EMERGENCY

### Description

This message occurs when a HBA port is taken offline due to read-after-write verification errors. This AutoSupport message and associated EMS messages are issued once for a given port. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Hardware components might need to be replaced. Contact NetApp technical support for assistance.

### Syslog Message

Call home for %s. Keep the adapter offline and contact NetApp technical support for assistance.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**adapterName** (STRING): Name of the adapter.

# callhome.hm events

## callhome.hm.alert.critical

### Severity

ALERT

### Description

This message occurs when an error is detected during the periodic health health monitoring of system health monitors. System health monitors create alerts for potential problems detected while monitoring the subsystem. The alerts contain information about probable cause along with recommended actions to rectify the problem. This message is generated after every polling interval if new alerts are detected during this interval. Polling interval is configurable. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run the "system health alert show" command to view details about new alerts in the system. Details of the alerts can be found in the AutoSupport. To correct the alert conditions, refer to the corrective actions described in the alert. Contact NetApp technical support for additional information.

### Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.hm.alert.major

### Severity

ALERT

### Description

This message occurs when an error is detected during the periodic health health monitoring of system health monitors. System health monitors create alerts for potential problems detected while monitoring the subsystem. The alerts contain information about probable cause along with recommended actions to rectify the problem. This message is generated after every polling interval if new alerts are detected during this interval. Polling interval is configurable. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run the "system health alert show" command to view details about new alerts in the system. Details of the alerts can be found in the AutoSupport. To correct the alert conditions, refer to the corrective actions described in the alert. Contact NetApp technical support for additional information.



## Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.hm.alert.minor

### Severity

ERROR

### Description

This message occurs when an error is detected during the periodic health monitoring of system health monitors. System health monitors create alerts for potential problems detected while monitoring the subsystem. The alerts contain information about probable cause along with recommended actions to rectify the problem. This message is generated after every polling interval if new alerts are detected during this interval. Polling interval is configurable. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run the "system health alert show" command to view details about new alerts in the system. Details of the alerts can be found in the AutoSupport. To correct the alert conditions, refer to the corrective actions described in the alert. Contact NetApp technical support for additional information.

## Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.hostport events

### callhome.hostport.bypass

### Severity

ALERT

### Description

This message occurs when an Embedded Switched Hub (ESH) bypasses a host port due to excessive errors on the port. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Depending on your system configuration, you might need to replace the ESH Module or replace the Fibre Channel cable to the host port. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for HOST\_PORT\_FAIL - Bypassed by ESH

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.hostport.disable

## Severity

ALERT

## Description

This message occurs when an I/O Module (IOM) disables a host port due to excessive errors on the port. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Depending on your system configuration, you might need to replace the IOM or ensure that a good cable is attached to the host port. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for HOST\_PORT\_FAIL - Disabled by IOM

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.hwaccel events

### callhome.hwaccel.reset

## Severity

NOTICE

## Description

This message occurs when the hardware accelerator is being reset and collects related diagnostic data concerning the device. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.ibretimerprog events

## callhome.ibretimerprog.fail

### Severity

EMERGENCY

### Description

This message occurs when InfiniBand (IB) retimer programming of the high-availability (HA) interconnect link failed after multiple retry attempts. The affected InfiniBand HA interconnect link is down. The failure can be due to any of the following reasons: InfiniBand device not found, GPIO read failure, QSFP register dump failure, QSFP parse failure, or retimer failure. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

This failure can be due to defective hardware or cabling. Contact NetApp technical support for further assistance.

### Syslog Message

Call home for INFINIBAND RETIMER PROGRAMMING FAILURE

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.invoke events

## callhome.invoke

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport® message, and the "type" of that message is set to "test". This message contains very little content and is intended for testing AutoSupport® end-to-end. When this message arrives at NetApp technical support, a confirmation email is sent to email destinations specified by the -to parameter. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this ASUP. If this is set, this ASUP is delivered to this alternate

destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If zero then obey the normal "state" and "support" flags. If this is "one", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.invoke.all

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport® message, and the "type" of that message is set to "all". For the benefit of technical support, the message also contains troubleshooting content, which might lengthen the collection time required. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

User triggered complete call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this ASUP. If this is set, this ASUP is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If zero then obey the normal "state" and "support" flags. If this is "one", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.invoke.remotesplog

### Severity

NOTICE

### Description

This message occurs when an administrator invokes an AutoSupport message on a node to collect the logs from the service processor of another node in the cluster. If the system is configured to do so, it generates and transmits an AutoSupport message (or 'call home' message) to NetApp technical support and to the configured destinations.

### Corrective Action

(None).

### Syslog Message

Call home (Remote SP LOG) for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this ASUP. If this is set, this ASUP is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If zero then obey the normal "state" and "support" flags. If this is "one", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.invoke.splog

### Severity

NOTICE

### Description

This message occurs when an administrator invokes an AutoSupport message for Service Processor (SP) log collection from ONTAP® CLI. If your system is configured to do so, it generates and transmits an AutoSupport message (or 'call home' message) to NetApp technical support and to the configured destinations.

### Corrective Action

(None).

### Syslog Message

Call home (SP LOG %s)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this ASUP. If this is set, this ASUP is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If zero then obey the normal "state" and "support" flags. If this is "one", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**for\_node\_id** (STRING): Node where the log is being collected (local or partner node).

## callhome.invoke.splog.full

### Severity

NOTICE

### Description

This message occurs when an administrator invokes an AutoSupport message for full Service Processor (SP) or baseboard management controller (BMC) log collection from ONTAP® CLI. If your system is configured to do so, it generates and transmits an AutoSupport message (or 'call home' message) to NetApp technical support and to the configured destinations.

### Corrective Action

(None).

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.iscsi events

## callhome.iscsi.hba.crash

### Severity

ALERT

### Description

This message occurs when the iSCSI host bus adapter (HBA) crashes. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If the HBA is in use, reboot your system to reset the HBA and retrieve the crash-dump file. Contact NetApp technical support and provide the crash-dump file for support purposes.

## Syslog Message

Call home for ISCSI HBA CRASH

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.isp events

## callhome.isp.adapter.bad

### Deprecated

Deprecated as of ONTAP 9.2. Last used in R8.2.x.

### Severity

ERROR

### Description

This message occurs when the system detects multiple Static Random Access Memory (SRAM) parity errors within a 30 day period. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Replace the Fibre Channel adapter that is exhibiting subsequent parity errors to insure continued operations. Look for the "fci.adapter.parity:info" EMS message in the event log to identify the adapter that is encountering the SRAM parity error and needs to be replaced. You can also check the output of the "FC-

STATS" section in AutoSupport and look for the value of link state counter "adap\_sram\_parity\_error\_cnt" which will be more than one for the concerned adapter. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for SRAM PARITY ERRORS ON FIBRE CHANNEL ADAPTER - replace adapter

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.jpc events**

### **callhome.jpc.selection.fail**

#### **Severity**

ALERT

#### **Description**

This message occurs when the Kernel Cluster Services (KCS) / Cluster Transaction Manager (CTRAN) detects that the algorithm to reselect the Join Proposal Coordinator (JPC) is failing repeatedly. The JPC is reselected when the node that hosted the JPC goes out of quorum. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### **Corrective Action**

Call NetApp technical support immediately for further assistance.

#### **Syslog Message**

Call home for JOIN PROPOSAL COORDINATOR SELECTION FAILING.

#### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## **callhome.kernel events**

### **callhome.kernel.core**

#### **Deprecated**

Deprecated as of ONTAP 9.3. This event is replaced by the callhome.kernel.coreinfo event.

#### **Severity**

NOTICE

#### **Description**

This message occurs when you invoke an AutoSupport for kernel coredump files. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Kernel coredump callhome for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If "0" then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this coredump AutoSupport message.

**core\_filename** (STRING): Coredump filename to be included in the coredump AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.kernel.coreinfo

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport for kernel coredump files. It contains content relevant to userspace cores. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Kernel coredump callhome for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If "0" then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this coredump AutoSupport message.

**core\_filename** (STRING): Coredump filename to be included in the coredump AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID



# callhome.link events

## callhome.link.change.stuck

### Severity

ALERT

### Description

This message occurs when a network port link state change cannot be confirmed for more than 5 minutes because another configuration task is still in progress. Network port states can change when the node reboots or is administratively changed by a system administrator. The affected port might still be operating in the state before the attempted change. While this condition exists, additional configuration changes or status queries will not be possible. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The link state change request will keep retrying until it succeeds. If the request does not succeed, it indicates that there is a configuration task that has not correctly completed and the node should be rebooted. For further assistance, contact NetApp technical support.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.management events

## callhome.management.log

### Severity

NOTICE

### Description

This message occurs when the MANAGEMENT\_LOG AutoSupport or 'call home' event is sent. This message occurs on a daily basis and is sent from clustered nodes only. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the node. If your node is configured to do so it generates and transmits an AutoSupport (or 'call home') to NetApp technical support and to the configured destinations on a daily basis. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for MANAGEMENT\_LOG

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.management.log.rotate

### Severity

NOTICE

### Description

This message occurs when the (MANAGEMENT\_LOG ROTATE) AutoSupport or 'call home' event is sent. This callhome message is a supplement to the daily callhome.management.log message. The callhome.management.log.rotate message occurs whenever a log file rotates due to size. This message avoids log information gaps when log files rotate. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the node. If your node is configured to do so it generates and transmits an AutoSupport (or 'call home') to NetApp technical support and to the configured destinations on a daily basis. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home %s for filename: %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**filename** (STRING): The name of the log file that rotated.

## callhome.mcc events

### callhome.mcc.auso.2min

### Severity

ALERT

### Description

This message occurs when a MetroCluster(tm) automatic unplanned switchover operation has completed and the outage window was longer than 120 sec. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check why switchover was triggered by using the "event log show" command and search for event "mcc.auso.triggered". Use the "metrocluster operation show" command to view the details of the switchover operation. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multinode ASUP.

## callhome.mcc.auso.complete

### Severity

ALERT

### Description

This message occurs when a MetroCluster(tm) automatic unplanned switchover operation has completed. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check why switchover was triggered by using the "event log show" command and search for event "mcc.auso.triggered". Use the "metrocluster operation show" command to view the details of the switchover operation. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multinode ASUP.

## callhome.mcc.auso.trig.fail

### Severity

EMERGENCY

### Description

This message occurs when a MetroCluster(tm) automatic unplanned switchover trigger failed to initiate a switchover operation. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Verify the health of each node and that automatic unplanned switchover is enabled on both by using the "metrocluster node show" command. Use the "metrocluster switchover -forced-on-disaster true" command to switch over the DR-site storage manually. For further assistance, contact NetApp technical support.

## Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.mcc.auso.trig.shut**

### **Severity**

ALERT

### **Description**

This message occurs when automatic unplanned switchover (AUSO) in a MetroCluster(tm) IP configuration is triggered due to remote site shutdown.

### **Corrective Action**

Check the hardware conditions at the remote site for the cause of the environmental shutdown. Use the "metrocluster operation show" command to view the details the switchover operation. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Callhome for %s.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.mcc.clst.dm.wd.gvup**

### **Severity**

ALERT

### **Description**

This message occurs when the Cluster DM Watchdog thread has given up after retrying for over an hour. Cluster configuration replication services are affected and configuration updates might be lost. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Examine the event management system (EMS) logs for any related issues. Perform any provided corrective actions. For further assistance, contact NetApp technical support.

### **Syslog Message**

Call home for METROCLUSTER CLUSTER DM WATCHDOG GAVE UP

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.mcc.cluster\_dm.stream\_in\_error**

### **Severity**

ALERT

### **Description**

This message occurs when cluster-domain sender/receiver stream has an error. The stream type is indicated through the flag. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for METROCLUSTER CLUSTER\_DM SENDER/RECEIVER STREAM\_IN\_ERROR

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**stream** (STRING): Stream is sender or receiver.

**callhome.mcc.config.complete****Severity**

NOTICE

**Description**

This message occurs when the "metrocluster configure" operation is successful. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for METROCLUSTER CONFIGURE COMPLETE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.mcc.fcvi.reset****Severity**

NOTICE

**Description**

This message occurs when fcvi adapter has been reset. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for METROCLUSTER FCVI RESET

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mcc.heal\_aggr.failed

### Severity

ALERT

### Description

This message occurs when metrocluster heal aggregates operation has failed. The failure could be partial or complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run 'metrocluster operation show' to determine the failure reason and take appropriate corrective measure. Retry the operation following the corrective measure. If the failures persists or need assistance, contact NetApp technical support

### Syslog Message

Call home for METROCLUSTER HEAL\_AGGR FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.heal\_root.failed

### Severity

ALERT

### Description

This message occurs when metrocluster heal root-aggregates operation has failed. The failure could be partial or complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run 'metrocluster operation show' to determine the failure reason and take appropriate corrective measure. Retry the operation following the corrective measure. If the failures persists or need assistance, contact NetApp technical support

### Syslog Message

Call home for METROCLUSTER HEAL\_ROOT FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.network.stats

### Severity

NOTICE

## Description

This message occurs when a MetroCluster over IP configuration detects network behavior change over the inter-cluster links (ISL) that connect sites. Diagnostic information is collected as part of this event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for %s. Number of times network changed since last callhome: %d.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**event\_count** (INT): Number of times network behavior change was flagged by the analysis engine since the last callhome message was sent.

## callhome.mcc.okm.sync.needed

### Severity

ALERT

## Description

This message occurs when clusters in a MetroCluster(tm) configuration have inconsistent Onboard Key Manager (OKM) data, including differing cluster passphrases or differing OKM hierarchies. The "security key-manager onboard sync" command should be run to prevent any failures in future switchover or switchback events. This message will be generated on both the local cluster and the peer cluster. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for METROCLUSTER OKM SYNC NEEDED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mcc.personality

### Severity

ALERT

## Description

This message occurs when the All-Flash Optimized node in a MetroCluster configuration detects that it has a different personality setting. This error can occur when a node loses its All-Flash Optimized setting. When this happens, the node may not function properly. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination

and resolution.

### Corrective Action

A node configured with All-Flash Optimized personality recognizes only SSD disks and ignores hard disk drives(HDD). The personality is determined at boot time by checking the bootarg value "bootarg.init.flash\_optimized". If the intended configuration of the node is All-Flash Optimized, any HDDs need to be removed and the bootarg needs to be set to true using the "setenv bootarg.init.flash\_optimized true" command in maintenance mode. Otherwise, use the "unsetenv bootarg.init.flash\_optimized" command to support HDDs on the system.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mcc.sb.vetoed

### Severity

NOTICE

### Description

This message occurs when the system detects that a MetroCluster (tm) switchback operation has been vetoed, and creates an AutoSupport message for delivery to NetApp technical support for further analysis.

### Corrective Action

Enter "metrocluster operation show" to determine the failure reason and take appropriate corrective measures as indicated by the output of this command. Retry the operation following the corrective measure. If the failures persist, or you need further assistance, contact NetApp technical support

### Syslog Message

Call home for METROCLUSTER SWITCHBACK VETOED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Correlating tag value used in the AutoSupport message. Other cluster nodes that have detected related events may also have sent AutoSupport messages with this value to NetApp technical support.

## callhome.mcc.sbca

### Severity

NOTICE

### Description

This message occurs when metrocluster switchback continuation agent (SBCA) is run on the disaster site. This happens as part of the metrocluster switchback operation run on the surviving site. This ASUP provides troubleshooting information from the site which is being recovered (disaster hit). If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.



**Corrective Action**

(None).

**Syslog Message**

Call home for METROCLUSTER SWITCHBACK CONTINUATION AGENT

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

**callhome.mcc.so.vetoed****Severity**

NOTICE

**Description**

This message occurs when the system detects that a MetroCluster (tm) switchover operation has been vetoed, and creates an AutoSupport message for delivery to NetApp technical support for further analysis.

**Corrective Action**

Enter "metrocluster operation show" to determine the failure reason and take appropriate corrective measures as indicated by the output of this command. Retry the operation following the corrective measure. If the failures persist, or you need further assistance, contact NetApp technical support

**Syslog Message**

Call home for METROCLUSTER SWITCHOVER VETOED

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Correlating tag value used in the AutoSupport message. Other cluster nodes that have detected related events may also have sent AutoSupport messages with this value to NetApp technical support.

**callhome.mcc.switchback.complete****Severity**

NOTICE

**Description**

This message occurs when metrocluster switchback operation is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for METROCLUSTER SWITCHBACK COMPLETE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.switchback.failed

### Severity

EMERGENCY

### Description

This message occurs when metrocluster switchback operation has failed. The failure could be partial or complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run 'metrocluster operation show' to determine the failure reason and take appropriate corrective measure. Retry the operation following the corrective measure. If the failures persists or need assistance, contact NetApp technical support

### Syslog Message

Call home for METROCLUSTER SWITCHBACK FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.switchover.2min

### Severity

NOTICE

### Description

This message occurs when MetroCluster(tm) switchover operation is complete and the client outage window has exceeded two minutes. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.switchover.complete

### Severity

NOTICE

### Description

This message occurs when metrocluster switchover operation is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for METROCLUSTER SWITCHOVER COMPLETE

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mcc.switchover.failed

### Severity

EMERGENCY

### Description

This message occurs when metrocluster switchover operation has failed. The failure could be partial or complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run 'metrocluster operation show' to determine the failure reason and take appropriate corrective measure. Retry the operation following the corrective measure. If the failures persists or need assistance, contact NetApp technical support

### Syslog Message

Call home for METROCLUSTER SWITCHOVER FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.mccip.events

### callhome.mccip.auso.disabled

### Severity

ALERT

## Description

This message occurs when automatic unplanned switchover (AUSO) in a MetroCluster(tm) IP configuration is disabled because the mediator mailbox disk is not healthy. If a site failure occurs, a manual switchover will be required to restore data access.

## Corrective Action

Use the "storage failover show -fields local-mailbox-disks" command to verify that all nodes in the cluster can access their respective mailbox disks. Use the (privilege: advanced) "storage iscsi-initiator show" command to verify that iSCSI connections to the mediator are up and healthy.

## Syslog Message

Callhome for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mccip.collectstats

### Severity

NOTICE

### Description

This message occurs when MetroCluster IP statistics are collected daily. TCP/IP networking statistics for diagnostic purposes are collected.

### Corrective Action

(None).

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mdb events

### callhome.mdb.recovery.unsuccessful

### Severity

EMERGENCY

### Description

This message occurs after an unsuccessful recovery of an inconsistent management database. The message is from the management processes identified by name in the auto-support subject line. The unsuccessful recovery results in the loss of all records contained in the 'job history' and 'rdb statistics' databases. Also, no additional records will be collected until the problem is resolved by NetApp technical support. The affected management process will function normally without these databases.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Call home for %s WARNING.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.mhost events

## callhome.mhost.core

### Deprecated

Deprecated as of ONTAP 9.3. This event is replaced by the callhome.mhost.coreinfo event.

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport for userspace coredump files. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

## Syslog Message

User space coredump callhome for %s

## Parameters

**subject** (STRING): AutoSupport message subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport message should be generated and delivered even if the AutoSupport configuration is disabled. If "0" then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this coredump AutoSupport message.

**core\_filename** (STRING): Coredump filename to be included in the coredump AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport message via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.mhost.coreinfo

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport for userspace coredump files. It contains content relevant to userspace cores. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport

message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

User space coredump callhome for %s

### Parameters

**subject** (STRING): AutoSupport message subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport message should be generated and delivered even if the AutoSupport configuration is disabled. If "0" then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this coredump AutoSupport message.

**core\_filename** (STRING): Coredump filename to be included in the coredump AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport message via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

## callhome.micro events

### callhome.micro.core

#### Severity

NOTICE

#### Description

This message occurs when a "micro-core" file is available. A "micro-core" file is a very small and simple binary image representing some state of ONTAP® and can be collected and reported to support.

Transmission of a normal core file which can be several gigabytes or larger can take a significant time so isn't done. The subject line provides the filename of the micro- core file which is included as part of a special AutoSupport normally only sent to support. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Call home for MICRO-CORE: %s

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.mlm events

## callhome.mlm.array.portOneController

### Severity

ALERT

### Description

This message occurs when all paths for an array LUN are connected to the same controller on the array. This configuration results in a single point of failure at the array. If an active-passive array is connected at the back end, then all the partner controller's paths go to the alternate controller, which results in a tug of war on mailbox I/O or other events that talk to the partner's array LUNs. If the system is running in Cluster-Mode, use the 'storage errors show' command to detect this error. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Reconfigure your back-end SAN so that array LUNs are presented from controllers in different fault domains on the target array. If you need further assistance, contact NetApp technical support.

### Syslog Message

Call home for ALL PATHS FOR THE ARRAY LUN ARE CONNECTED TO THE SAME CONTROLLER.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.mlm.lun.OnePath

### Severity

ALERT

### Description

This message occurs when the logical devices (LDEVs) on an array LUN or group of array LUNs are connected using only one path. A proper configuration requires the LDEVs to be connected to at least two paths; otherwise, there is the potential for data loss caused by a single path failure. If an array LUN is assigned, an EMS message is logged to the console at least once a day indicating that it is a single-path assigned array LUN. If the system is running in Cluster-Mode, use the 'storage errors show' command to identify the error. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Verify the configuration and identify the LDEVs that are connected to only one path. After you identify the LDEVs, change the system configuration to ensure that there are at least two paths to all these LDEVs. If you need further assistance, contact NetApp technical support.

### Syslog Message

Call home for ARRAY LUN CONNECTED USING ONLY ONE PATH.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.multi events

## callhome.multi.dsk.cksum.err

### Severity

EMERGENCY

### Description

This message occurs when checksum verification errors are seen on two or more disks connected to the same HBA port, and the port is not taken offline or the "disk.offline\_adapter.enable" option is disabled. This AutoSupport® message and associated EMS messages are issued once for a given port. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for assistance with troubleshooting the reason for the write verification errors. Hardware components might need to be replaced.

### Syslog Message

Call home for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**adapterName** (STRING): Name of the adapter.

# callhome.ndo events

## callhome.ndo.cr.canceled

### Severity

NOTICE

### Description

This message occurs when the controller replacement operation has been canceled by the administrator. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

The controller replacement operation on the cluster has been canceled. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

### Syslog Message

(None).



**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ndo.cr.complete****Severity**

NOTICE

**Description**

This message occurs when the controller replacement operation has been completed successfully in the HA group. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for CONTROLLER REPLACEMENT COMPLETE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ndo.cr.dryrun.done****Severity**

NOTICE

**Description**

This message occurs when the dry run of the controller replacement operation has completed for the HA pair. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for CONTROLLER REPLACEMENT DRY RUN COMPLETE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ndo.cr.nso.canceled****Severity**

NOTICE

**Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been canceled. If your system is configured to do so, it generates and

transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

The controller replacement NSO operation on the cluster has been canceled. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

### **Syslog Message**

(None).

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ndo.cr.nso.complete**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been completed successfully. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

(None).

### **Syslog Message**

Call home for CONTROLLER REPLACEMENT NSO COMPLETE

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ndo.cr.nso.mccip.canceled**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been canceled. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

The controller replacement NSO-MCCIP operation on the cluster has been canceled. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ndo.cr.nso.mccip.complete

## Severity

NOTICE

## Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been completed successfully. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for CONTROLLER REPLACEMENT NSO-MCCIP COMPLETE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ndo.cr.nso.mccip.paused.err

## Severity

ALERT

## Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been paused due to an error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

The controller replacement (NSO-MCCIP) operation of the cluster has been paused due to an error. Determine the cause of the error by using the "system controller replace show" command. Correct the cause of the error, and then issue the "system controller replace resume" command to continue the controller replacement of the cluster.

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ndo.cr.nso.mccip.started**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been started. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

The controller replacement NSO-MCCIP operation of the cluster has started. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

### **Syslog Message**

(None).

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ndo.cr.nso.paused.err**

### **Severity**

ALERT

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been paused due to an error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

The controller replacement (NSO) operation of the cluster has been paused due to an error. Determine the cause of the error by using the "system controller replace show" command. Correct the cause of the error, and then issue the "system controller replace resume" command to continue the controller replacement of the cluster.

### **Syslog Message**

(None).

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ndo.cr.nso.started**

### **Severity**

NOTICE

## Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been started. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

The controller replacement NSO operation of the cluster has started. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ndo.cr.paused.err

### Severity

ALERT

## Description

This message occurs when the controller replacement operation has been paused due to an error in the HA group. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

The controller replacement operation of the cluster has been paused due to an error. Determine the cause of the error by using the "system controller replace show" command. Correct the cause of the error, and then issue the "system controller replace resume" command to continue the controller replacement of the cluster.

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ndo.cr.started

### Severity

NOTICE

## Description

This message occurs when the controller replacement operation has started in the HA group. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

The controller replacement operation of the cluster has started. Issue the commands "system controller replace show" and "system controller replace show-details" to see the status of the operation."

### **Syslog Message**

(None).

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.netif events**

### **callhome.netif.fatal.error**

#### **Severity**

ERROR

#### **Description**

This message occurs when Network Interface Card encountered a fatal error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### **Corrective Action**

Check for the following EMS message in the EMS log that has the EMS event name "netif.fatal.err". For example: "...netif.fatal.err:crit]: The network device in slot 1 Encountered Fatal Error" The netif.fatal.err EMS message provides the slot number of the NIC device which encountered fatal error. If you need assistance, contact NetApp technical support

#### **Syslog Message**

Call home for NETWORK INTERFACE CARD FATAL ERROR.

#### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

### **callhome.netif.init.failed**

#### **Severity**

ALERT

#### **Description**

This message occurs when initialization of a network interface failed due to an unexpected software error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### **Corrective Action**

Review the cabling and network configuration. Collect faulty NIC statistics. If you need assistance, contact NetApp technical support

## Syslog Message

Call home for NETWORK INTERFACE INITIALIZATION FAILURE, port %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.netif.uncorEccError

## Severity

ALERT

## Description

This message occurs when a network interface controller encountered an unrecoverable ECC error. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Review the cabling and network configuration. Replace the faulty NIC device. If you need assistance, contact NetApp technical support

## Syslog Message

Call home for NETWORK INTERFACE UNRECOVERABLE ECC ERROR, port %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.netinet events

### callhome.netinet.dup.clustIP

## Severity

EMERGENCY

## Description

This message occurs when a new node that is trying to join the cluster is using the same IP address that is configured on this node's cluster interface. This can prevent data from being served to clients and potentially bring down the entire cluster. This event generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Use the Ethernet (MAC) address with the "network port show -mac <mac>" command to find network port. Use "network interface show -curr-port <port>" command to identify the cluster interface that is configured with the duplicate IP address. Remove the LIF on the cluster interface of the new node in the cluster network by using the **net int delete -vserver <vserver\_name > -lif <lif\_name** command. After removing the LIF, assign a new, unique IP address to the cluster interface.

## Syslog Message

Call home for DUPLICATE CLUSTER IP.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**ipaddr** (STRING): Duplicate IP address that was detected on the cluster network.

**ethaddr** (STRING): Ethernet (MAC) address that uniquely identifies the node configured with the duplicate IP address.

# callhome.nht events

## callhome.nht.data

### Severity

NOTICE

### Description

This message occurs when the weekly AutoSupport or 'call home' event is sent for any drive NHT data. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the system. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for DRIVE NHT DATA

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.no events

## callhome.no.inodes

### Severity

ERROR

### Description

This message occurs when a volume is out of inodes, which refer to individual files, other types of files, and directories. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Modify the volume's maxfiles (maximum number of files) to increase the inodes on the affected volume. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for OUT OF INODES



## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.nse events

## callhome.nse.ak.check.failed

### Severity

EMERGENCY

### Description

This message occurs when the node's key-management database lacks authentication keys (AKs) needed by self-encrypting drives that are owned by the node. The affected drives can be FIPS-certified drives (variously called FIPS drives, NetApp Storage Encryption drives, and NSE drives) or non-FIPS-certified self-encrypting drives (SEDs). A FIPS drive or SED is at risk of permanent data loss if authentication fails repeatedly after it is power-cycled. When FIPS keys are missing, operations such as encrypting sanitize and firmware download might also not succeed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Display the available keyIDs for encrypting drives using the "security key-manager key query -key-type NSE-AK" command; ONTAP uses these values for all its encrypting drives, whether FIPS/NSE drives or SEDs. Display keyIDs that the drives need using the command "storage encryption disk show", and for FIPS drives, "storage encryption disk show -fips". Each keyID in the disk lists must be present in the key-manager output, except for the default "0x0" (MSID) and "" (null) keyIDs that are not maintained by the key management system. Contact NetApp technical support for assistance. Avoid drive power-cycles and AK modifications until the issue is corrected.

### Syslog Message

Call home for %s, disk "%s".

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**disk** (STRING): Name of the first affected drive detected.

# callhome.nvdimm events

## callhome.nvdimm.cecc

### Severity

ERROR

### Description

This event occurs when excessive correctable error-correcting code (ECC) memory errors are reported by the NVDIMM. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Use the "system environment sensors show -hidden true" clustershell command to view the NVDIMM sensors values. Contact NetApp technical support for further assistance.

### **Syslog Message**

Call home for %s.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.nvdimm.failure**

### **Severity**

EMERGENCY

### **Description**

This event occurs when the NVDIMM card has experienced a critical hardware failure. The system may shut down to avoid potential data loss. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Use the following clustershell commands to get more information about the failure: 1. "event log show -node localhost -event nvdimm\*" 2. "system fru-check show -node localhost -ftype nvdimm -instance" Contact NetApp technical support to replace the NVDIMM module.

### **Syslog Message**

Call home for %s.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.nve events**

### **callhome.nve.digest.mismatch**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the digest of an NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) key being restored does not match the expected key digest. Because the key being restored is not the same key expected by the system, the volume associated with the NVE or NAE key will not be able to be onlined, resulting in data loss. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Call home for %s. key ID: %s, Key Server: %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**keyId** (STRING): Key ID of the key being restored that had the digest mismatch.

**keyServer** (STRING): Address or name of the key server from which the key is being restored.

## callhome.nvme events

### callhome.nvme.exhausted

#### Severity

ALERT

#### Description

This message occurs when a directly attached NVMe module has worn out and is at the end of its write life. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Contact technical support for assistance with possible replacement of the module.

#### Syslog Message

Callhome for %s.

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**serial** (STRING): Serial number of the NVMe device.

**model** (STRING): Model number of the NVMe device.

**slot\_subslot** (STRING): PCI slot/subslot of the NVMe device.

### callhome.nvme.offline

#### Severity

ALERT

#### Description

This message occurs when a directly attached NVMe module goes offline. When a module goes offline, it has either had a critical failure or is no longer reliable. In either case, the card is no longer functioning properly and was taken offline. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact technical support for assistance with possible replacement of the module.

### Syslog Message

Callhome for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**serial** (STRING): Serial number of the NVMe device.

**model** (STRING): Model number of the NVMe device.

**slot\_subslot** (STRING): PCI slot/subslot of the NVMe device.

**code** (LONGINTEX): Bit-encoded error that caused the event.

## callhome.nvmem events

### callhome.nvmem.volt.hi

#### Severity

ALERT

#### Description

This message occurs when the sensor for the Non-Volatile memory (NVMEM) card lithium battery has an excessively high voltage. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

First, correct any environmental or battery problems. If the problem continues, replace the controller module. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for NVMEM VOLTAGE HIGH

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.nvmem.volt.lo

#### Severity

ALERT

#### Description

This message occurs when the sensor for the Non-Volatile memory (NVMEM) card lithium battery has an excessively low voltage. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

First, correct any environmental or battery problems. If the problem continues, replace the controller module. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for NVMEM VOLTAGE LOW

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.nvram events

## callhome.nvram.cecc

### Severity

ERROR

### Description

This message occurs when excessive correctable error-correcting code (ECC) memory errors are reported by the NVRAM complex. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "system environment sensors show -hidden true" cluster shell command to view the current "NV Correctable ECC count" and "NV Correctable ECC rate" values. Contact NetApp technical support for further assistance.

## Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.nvram.failure

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM card has experienced a critical hardware failure. The system will shut down to avoid potential data loss. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support to replace the NVRAM card.

## Syslog Message

Call home for NVRAM FAILURE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.nvram.iofs.fail

### Severity

ERROR

### Description

This message occurs in a FireStorm instance when a read or write to its nonvolatile memory (NVRAM) device fails. The instance will panic. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Restart the FireStorm instance. If the issue persists, contact NetApp technical support.

### Syslog Message

Call home for NVRAM failure at %s, I/O type %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**io\_type** (STRING): Type of failed operation (read or write).

**tag** (STRING): Supplemental event identifier.

**error** (LONGINT): I/O error code.

## callhome.nvram.openfs.fail

### Severity

ERROR

### Description

This message occurs when a FireStorm instance is unable to open its nonvolatile memory (NVRAM) device. The instance will panic.

### Corrective Action

Restart the instance. If the issue persists, contact NetApp technical support.

### Syslog Message

Call home for open failure at %s on NVRAM file %s, error %d.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**nvram\_backing** (STRING): Name of the NVRAM backing file.

**fileio\_error** (LONGINT): Open error code.

## callhome.okm events

### callhome.okm.sync.needed

### Severity

ALERT

## Description

This message occurs when the cluster has reached an inconsistent state with respect to the onboard key hierarchy. This is an alert that the "security key-manager onboard sync" command should be run on the local cluster to prevent any future usage failures with onboard key management. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to Netapp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for CLUSTER OKM SYNC NEEDED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.ondemand events

## callhome.ondemand.invoke

## Severity

NOTICE

## Description

This message normally occurs when NetApp technical support has remotely requested an AutoSupport or 'call home' message of type TEST. This request is accomplished by use of the AutoSupport OnDemand Server (AODS) service. The intent of this AutoSupport is to collect only the information necessary to understand the configuration of the system and quickly deliver it. For this reason, the data collection is bounded by time taken to collect data and quantity of data collected. With less data collected, delivery over the internet is expected to be faster. The overall process from remote activation to delivery should be measurable in minutes. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport. If set, this AutoSupport is delivered to this alternate destination.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

## callhome.ondemand.invoke.all

### Severity

NOTICE

### Description

This message normally occurs when NetApp technical support has remotely requested an AutoSupport or 'call home' message of type ALL. This request is accomplished by use of the AutoSupport OnDemand Server (AODS) service. The intent of this AutoSupport is to collect more data to aid in problem resolution. The collection is bounded by time and size limits but these are much larger as compared to most other AutoSupport messages. Collecting a larger amount of AutoSupport data can help support personnel diagnose difficult problems. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport. If set, this AutoSupport is delivered to this alternate destination.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

## callhome.ondemand.perf.snap

### Severity

NOTICE

### Description

This message normally occurs when NetApp technical support has remotely requested an AutoSupport or 'call home' message of type PERFORMANCE. This request is accomplished by use of the AutoSupport OnDemand Server (AODS) service. The intent of this AutoSupport is to collect statistics and other information to aid in understanding operations and other usage analysis related to networking, storage, memory and processor. This analysis might result in recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the system. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

(None).



## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport. If set, this AutoSupport is delivered to this alternate destination.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

# callhome.partner events

## callhome.partner.down

### Severity

EMERGENCY

### Description

This message occurs when the storage failover partner is down and takeover cannot begin. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Call home for PARTNER DOWN, TAKEOVER IMPOSSIBLE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.performance events

## callhome.performance.archive

### Severity

NOTICE

### Description

This message occurs when you invoke an AutoSupport for performance archives. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Performance archive callhome for %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this AutoSupport message. If this is set, this AutoSupport message is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If 0 then obey the normal "state" and "support" flags. If this is "1", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**case\_number** (STRING): Case number to be associated with this performance archive AutoSupport message.

**temp\_tar\_path** (STRING): Tared performance archives to be included in this performance archive AutoSupport message.

**aod\_initiated** (INT): This flag indicates that NetApp technical support has remotely requested for this AutoSupport via AutoSupport OnDemand.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

**aod\_req\_id** (STRING): AOD request ID

## callhome.performance.data

### Severity

NOTICE

### Description

This message occurs when the weekly AutoSupport or 'call home' event is sent for performance data. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the system. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for PERFORMANCE DATA

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.performance.snap

### Severity

NOTICE

### Description

This message occurs some minutes after a reboot or storage takeover operation occurs so that performance information can be sent. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the system. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations.

Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for PERFORMANCE SNAPSHOT

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**URI** (STRING): Alternate destination for this ASUP. If this is set, this ASUP is delivered to this alternate destination.

**force** (INT): This flag indicates this AutoSupport should be generated and delivered even if the AutoSupport configuration is disabled. If zero then obey the normal "state" and "support" flags. If this is "one", then ignore the "state" and "support" configuration and deliver as if both "state" and "support" are enabled.

**sequence** (LONGINT): Sequence number to use to generate this Autosupport. Use value zero to auto-generate or a value higher than last used ASUP sequence number.

## callhome.pool events

### callhome.pool.aggr.noncomp

#### Severity

ALERT

#### Description

This message occurs when an aggregate does not comply with the current Capacity Pool entitlement, as determined by the Licensing Manager. In an HA configuration, noncompliant aggregates are prevented from being sent home during an HA takeover/giveback situation. In a non-HA configuration, these aggregates cannot be brought online if they are already offline. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Use the "system license license-manager check" command to determine whether the License Manager is accessible. Use the "system license show-aggregates -compliant false" command to find all aggregates that do not comply. From the web interface of the ONTAP® Select Deploy® instance, check the Capacity Pool licensing status, and ensure that sufficient capacity remains within the capacity pool used by this instance of ONTAP. If the capacity is insufficient, reduce the usage by instances of ONTAP that use the pool (for example, by deleting aggregates). If no reduction is possible, Contact NetApp technical support..

#### Syslog Message

Call home for %s.

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.raid events

## callhome.raid.adp.disabled

### Severity

NOTICE

### Description

This message occurs when auto-partitioning for root aggregate creation is disabled on a system. When auto-partitioning is disabled, the root aggregate is created on whole disks. Systems with a root aggregate on partitioned disks can devote more space to non-root aggregates than systems with a root aggregate on whole disks. If this message occurs after root aggregate creation, then further partitioning of the disks is disabled.

### Corrective Action

If the system is already initialized with root aggregate on partitioned disks, ONTAP will try to re-enable auto-partitioning after the system boots up and the root aggregate is online. The event `raid.autoPart.enabled` will be logged when auto-partitioning is re-enabled. To create a root aggregate on partitioned disks: 1. Halt the controller. 2. Replace any failed internal disks (see the "reason" field in the EMS message). 3. Reboot the controller. 4. Re-initialize the system. If this message is seen after the system initialization and the root aggregate is already created on partitioned disks, then the system will continue to function, as follows: 1. Expansion of existing aggregates created from partitioned disks may not work. 2. If a disk fails in a RAID group created out of partitioned disks, then a spare will fail to be partitioned and disk reconstruction will not start. 3. If a disk is marked for predictive failure in a RAID group created out of partitioned disks, then a spare will fail to be partitioned and the sick disk copy will not start. 4. If disks are available, then manually partition the disks and keep them as spares. 5. If you were previously directed to disable ADP by NetApp technical support, reach out to NetApp technical support to re-enable ADP.

### Syslog Message

Disk auto-partitioning is disabled on this system: %s.

### Parameters

**subject** (STRING): Reason that auto-partitioning was disabled.

## callhome.raid.no.recover

### Severity

EMERGENCY

### Description

This message occurs when too many errors have been detected by the RAID (Redundant Array of Inexpensive Disks) software protection and recovery is not possible. The detected errors could be related to checksum calculations, lost write detection and on disk medium errors in combinations with disk and NVRAM failures. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Immediately contact NetApp technical support for assistance to help diagnose the failures. If the diagnosis points to malfunctioning hardware components, these should be replaced.

## Syslog Message

Call home for RAID DETECTED UNRECOVERABLE ERRORS

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.raid.over.max

## Severity

ALERT

## Description

This message occurs when the maximum raw capacity of the file system (aggregate or RAID file system) is reached for this software release. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Call home for Exceeded maximum raw capacity

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.raidtree events

### callhome.raidtree.assim

## Severity

EMERGENCY

## Description

This message occurs when an aggregate experiences a configuration corruption error during assimilation of aggregate structures. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Call home for RAIDTREE ASSIMILATION FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.reboot events

## callhome.reboot

### Severity

NOTICE

### Description

This message occurs after an appliance reboot—the appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If reboots continue to occur, contact NetApp technical support.

### Syslog Message

Call home for REBOOT

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.agentwatchdog

### Severity

ALERT

### Description

This message occurs when the system experiences an internal hardware watchdog timer reset. The system reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If you continue to experience watchdog timer resets, contact NetApp technical support.

### Syslog Message

Call home for REBOOT (agent watchdog reset).

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.fw

### Severity

NOTICE

### Description

This message occurs when the appliance reboots after loading firmware—the appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an

AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for REBOOT (after entering firmware)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.reboot.giveback****Severity**

NOTICE

**Description**

This message occurs when the node boots to normal operations because of a successful storage giveback from the partner. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for REBOOT (after giveback)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.reboot.glitch****Severity**

NOTICE

**Description**

This message occurs when the appliance halts and reboots due to problems with inconsistent power. The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

If you continue to experience 'power glitch' reboots, contact NetApp technical support.

**Syslog Message**

Call home for REBOOT (power glitch)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.reboot.halt****Severity**

NOTICE

**Description**

This message occurs when the appliance reboots after you use the 'halt' command to shut it down. The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for REBOOT (halt command)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.reboot.ioxm.removal****Severity**

NOTICE

**Description**

This message occurs when the appliance halts and reboots due to the removal of the IOXM without powering down the appliance first. The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Do not remove the IOXM without powering down the appliance first.

**Syslog Message**

Call home for REBOOT (IOXM removal)

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.reboot.panic****Severity**

ALERT



## Description

This message occurs when the appliance reboots after a panic (also known as a 'system crash'). The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Call home for REBOOT (panic)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.poweron

### Severity

NOTICE

## Description

This message occurs when the appliance reboots when it is powered on. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you continue to experience 'power on' reboots, contact NetApp technical support.

## Syslog Message

Call home for REBOOT (power on)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.reboot

### Severity

NOTICE

## Description

This message occurs when you reboot the appliance with the 'reboot' command. The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for REBOOT (reboot command)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.sk.halt

## Severity

NOTICE

## Description

This message occurs when the system is halted automatically to prevent damage or possible loss of data. The system reports this event after the reboot is complete. An internal halt might be caused by one of the following reasons: 1. If a disk drive fails and there are no available spares, the system enters degraded mode. The appliance halts automatically to prevent a double disk drive failure, and possible loss of data, if it runs in degraded mode for the time interval given by the following command: "storage raid-options show -name raid.timeout -node local" The default timeout is usually 24 hours. If a spare disk drive becomes available while the system is running in degraded mode, the system immediately begins rebuilding the failed disk drive. 2. The system also halts automatically if the environmental conditions, such as temperature or power, are not within the safe operating limits or are inadequate to prevent possible damage. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you continue to experience 'internal halt' reboots, contact NetApp technical support.

## Syslog Message

Call home for REBOOT (internal halt)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.sk.reboot

## Severity

NOTICE

## Description

This message occurs when you halt, reset, or turn off power to the appliance without completing an orderly halt or shutdown sequence before rebooting. The appliance reports this event after the reboot is complete. To keep your data safe, always use the 'halt' command to halt the appliance, and then wait for an orderly halt or shutdown to be complete before you power down the appliance or perform any hardware maintenance. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you continue to experience 'internal reboot' reboots, contact NetApp technical support.

## Syslog Message

Call home for REBOOT (internal reboot)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.sp.rst

## Severity

NOTICE

## Description

This message occurs when the appliance reboots when it has been reset through the Service Processor (SP) 'system reset' command. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

(None).

## Syslog Message

Call home for REBOOT (sp system reset)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.takeover

## Severity

NOTICE

## Description

This message occurs when the partner appliance storage is taken over successfully by this appliance. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for PARTNER REBOOT (CONTROLLER TAKEOVER)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.takeover.p

## Severity

NOTICE

## Description

This message occurs when the partner appliance storage is taken over successfully by this appliance. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for PARTNER REBOOT (CONTROLLER TAKEOVER ON PANIC)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.unknown

### Severity

NOTICE

## Description

This message occurs after an appliance reboot—the appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If reboots continue to occur, contact NetApp technical support.

## Syslog Message

Call home for REBOOT

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.reboot.watchdog

### Severity

ALERT

## Description

This message occurs when the appliance experiences an internal hardware watchdog timer reset. The appliance reports this event after the reboot is complete. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If you continue to experience watchdog timer resets, contact NetApp technical support.

## Syslog Message

Call home for REBOOT (watchdog reset)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.rnode events

## callhome.rnode.splog

### Severity

NOTICE

### Description

This messages occurs when a new set of log files has been collected from a Service Processor (SP) of one or more non-operational nodes in the cluster. When an abnormal node reboot is detected by its SP, log files are created. When this node is operational again, the SP log files are collected and included in subsequent AutoSupport messages sent by the node. In cases where nodes cannot return to operation immediately, the SP's log files are collected by another node in the cluster and sent via this AutoSupport message. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

## Syslog Message

Call home for REMOTE NODE SP LOG %s %s %s %s %s %s %s %s %s %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**node1\_down\_fn** (STRING): Node-specific filename for the downfiler splog from the first node. If NULL is specified, this downfiler splog is not available.

**node1\_latest\_fn** (STRING): Node-specific filename for the latest splog from the first node.

**node2\_down\_fn** (STRING): Node-specific filename for the downfiler splog from the second node. If NULL is specified, this downfiler splog is not available.

**node2\_latest\_fn** (STRING): Node-specific filename for the latest splog from the second node.

**node3\_down\_fn** (STRING): Node-specific filename for the downfiler splog from the third node. If NULL is specified, this downfiler splog is not available.

**node3\_latest\_fn** (STRING): Node-specific filename for the latest splog from the third node.

**node4\_down\_fn** (STRING): Node-specific filename for the downfiler splog from the fourth node. If NULL is specified, this downfiler splog is not available.

**node4\_latest\_fn** (STRING): Node-specific filename for the latest splog from the fourth node.

**node5\_down\_fn** (STRING): Node-specific filename for the downfiler splog from the fifth node. If NULL is specified, this downfiler splog is not available.

**node5\_latest\_fn** (STRING): Node-specific filename for the latest splog from the fifth node.

# callhome.root events

## callhome.root.vol.recovery.reqd

### Severity

EMERGENCY

### Description

This message occurs when the root volume is not working properly and its recovery is required. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support to obtain the root volume recovery procedures and further assistance.

### Syslog Message

Call home for ROOT VOLUME NOT WORKING PROPERLY: RECOVERY REQUIRED.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.rtc events

### callhome.rtc.battery.low

### Severity

ERROR

### Description

This message occurs when the voltage for the Real-Time Clock (RTC) battery is low. The message is repeated once every hour until you replace the battery. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Replace the RTC battery as soon as possible. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for RTC BATTERY LOW

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sas events

### callhome.sas.init.failed

### Severity

ERROR

## Description

This message occurs when the SAS adapter fails to initialize the link. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp Global Services and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

If the adapter is in use, check cabling and connected equipment. If the adapter is connected to storage shelves, check connected I/O modules (IOM).

## Syslog Message

Callhome for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**adapterName** (STRING): Name of the associated SAS host bus adapter.

# callhome.sb events

## callhome.sb.pp.fcp.failo.err

### Severity

ERROR

### Description

This message occurs when the persistent port feature is not working during non-disruptive operations (NDO) such as takeover or giveback due to configuration errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sb.pp.fcp.init.err

### Severity

ERROR

### Description

This message occurs when persistent port object for a SAN FCP LIF fails to be configured during a LIF create operation. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sblade events

### callhome.sblade.import.susp

**Severity**

ERROR

**Description**

This message occurs when LUN import cannot resume after takeover/giveback. All LUN import sessions have been suspended. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Call home for LUN IMPORT SESSIONS SUSPENDED

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multinode AutoSupport.

### callhome.sblade.lif.init

**Severity**

NOTICE

**Description**

This message occurs during Vserver initialization when partially initialized volumes do not complete initialization within the specified maximum timeout value. On timer expiration the Vserver data LIFs are brought online regardless of the volume initialization status. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for SCSIBLADE LIF INIT TIMEOUT



**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sblade.lu.resync.to****Severity**

EMERGENCY

**Description**

This message occurs when the SCSI target internal resynchronization of a LUN times out. This LUN is unavailable until one or more nodes that could not complete the internal resynchronization are taken over and given back or rebooted. If your system is configured to do so, it generates and transmits an AutoSupport® (or "call home") message. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Call NetApp technical support for further assistance.

**Syslog Message**

Call home for SCSIBLADE LU RESYNC TIMEOUT

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sblade.lu.rst.hung****Severity**

EMERGENCY

**Description**

This message occurs when the SCSI target cannot complete an internal reset of a LUN. This LUN is unavailable until one or more nodes that could not complete the internal reset are taken over and given back or rebooted. If your system is configured to do so, it generates and transmits an AutoSupport® (or "call home") message. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Call NetApp technical support for further assistance.

**Syslog Message**

Call home for SCSIBLADE LU INTERNAL RESET HUNG

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sblade.mgmt.wedged****Severity**

EMERGENCY

## Description

This message occurs when the SAN management plane encounters an error that prevents further SAN management commands from being processed. This usually means that all the threads of one (or more) of the SAN management plane modules are stuck. SAN configuration changes on this node and other nodes in the cluster will fail with timeout errors until corrected. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Call NetApp technical support immediately for further assistance.

## Syslog Message

Call home for SCSIBLADE MGMT WEDGED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sblade.prop.fail

## Severity

EMERGENCY

## Description

This message occurs when an internal SAN management command fails. SAN management commands is not able to be completed for the Vserver which fails until the node is rebooted. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Refer to Knowledgebase article: 2017459 for recovery instructions.

## Syslog Message

Call home for SCSIBLADE PROPOSAL FAIL

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sblade.unavailable

## Severity

ERROR

## Description

This message occurs when the node has lost connectivity with the majority of the cluster and is unable to recover automatically. When this occurs both FCP and iSCSI traffic have been disabled on this node. Manual recovery is required to come back into service. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Reboot the node, and if the node is waiting for giveback then perform a giveback.

**Syslog Message**

Call home for SCSIBLADE SERVICE UNAVAILABLE

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sblade.vol.ev.lost****Severity**

ERROR

**Description**

This message occurs when LUNs (Logical Unit Numbers) in a volume are not accessible on all paths. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Call NetApp technical support for further assistance.

**Syslog Message**

Call home for SCSIBLADE VOLUME EVENT LOST

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

**callhome.sblade.vs.op.tt****Severity**

ERROR

**Description**

This message occurs when a management command or other internal operation takes too long for a LUN (Logical Unit Number). When this situation occurs, the corresponding Vserver cannot handle any other management changes for the LUN until this condition no longer persists. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Call NetApp technical support immediately for further assistance.

**Syslog Message**

Call home for SCSIBLADE VSERVER OPERATION TIMEOUT

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi node ASUP.

## callhome.scsi events

### callhome.scsi.tgt.misconfig

#### Severity

ERROR

#### Description

This message occurs when a LUN (Logical Unit Number) mapping or igroup type misconfiguration has been detected when the cluster interconnect has transitioned from down to up. The detected misconfiguration problems will typically include cases of having a LUN on each storage controller mapped to the same LUN for the same initiator. When the cluster interconnect is down, LUN mapping changes are prevented unless overridden. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Run the "lun config\_check" command which will list out the problems that caused this event and resolve this problem list. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for FILER SCSI TARGET MISCONFIGURED

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sdisk events

### callhome.sdisk.admin

#### Severity

NOTICE

#### Description

This message occurs when you fail a hot spare disk drive.

#### Corrective Action

(None).

#### Syslog Message

Call home for SPARE DISK ADMIN REMOVED

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.sdisk.fault

**Severity**

ERROR

**Description**

This message occurs when a spare disk drive fails. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sdisk.label****Severity**

ERROR

**Description**

This message occurs when a disk drive fails due to unexpected RAID metadata. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for SPARE DISK BAD LABEL

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sdisk.label.v****Severity**

ERROR

**Description**

This message occurs when a disk drive fails due to a version number in its RAID metadata that is inconsistent with the operating system version installed. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for SPARE DISK BAD LABEL VERSION

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sdisk.missing****Severity**

ERROR

**Description**

This message occurs when the system detects that a spare disk drive is not in a loop or is missing. This might also be due to a failed disk drive that is still in the chassis awaiting replacement. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate and replace the failed disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for SPARE DISK MISSING

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sdisk.noio****Severity**

ERROR

**Description**

This message occurs when disk drive device commands to a hot spare disk drive are failing with link errors, adapter errors, or timeouts. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for SPARE DISK NOT RESPONDING

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sdisk.shrank

### Severity

ERROR

### Description

This message occurs when a hot spare disk drive fails because its raw capacity was downsized due to a firmware change. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Locate the failed disk drive by looking for a disk drive whose red and amber LED fault lights are NOT lit. Replace the failed disk drive. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SPARE DISK RAWSIZE SHRANK

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ses events

### callhome.ses.drive.missing

### Severity

ERROR

### Description

This message occurs when the storage adapter detects a disk on the Fibre Channel loop or the SCSI bus, but enclosure services in the disk shelf does not report a disk with that address. This error indicates that the disk took an address other than what the disk shelf should have provided, or that enclosure services in the disk shelf cannot be contacted for address information. (Note: If the message "ses.drive.possShelfAddr" subsequently appears, follow the corrective actions in that message.) In this condition, the SCSI Enclosure Services (SES) process in the system might be unable to perform certain operations on the disk, such as setting failure LEDs or detecting disk swaps. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If the disk is not part of an active aggregate or volume, pull out the disk and push it back into the slot to ensure that the disk is in place. If the problem persists, insert a spare disk into the slot. If the problem is then gone, replace the original disk with the spare one. If the problem occurs to multiple disks on the same loop, check the I/O modules at the back of the disk shelves on that loop for errors.

### Syslog Message

Call home for SES: DRIVE MISSING

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**ChannelName** (STRING): Channel the disk shelf is on.

**LoopIds** (STRING): Loop or SCSI IDs of the missing disks.

## callhome.sfo events

### callhome.sfo.downrev

#### Severity

ERROR

#### Description

This message occurs when the appliance software boots and detects out-of-date motherboard firmware, which can cause serious issues. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Upgrade the motherboard open firmware (OFW) version to the recommended version for this software and appliance as soon as possible. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for HA GROUP DOWNREV BOOT FIRMWARE

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.sfo.giveback

#### Severity

NOTICE

#### Description

This message occurs when this appliance successfully completed a giveback of the storage to its partner appliance. The partner appliance is expected to be performing normal operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

(None).

#### Syslog Message

Call home for CONTROLLER GIVEBACK COMPLETE

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.sfo.miscount



**Severity**

ERROR

**Description**

This message occurs when each partner in the High Availability (HA) storage configuration detects a different number of disk drives than its partner. This error is typically transient, and can occur for the following reasons: 1) A disk drive failed. 2) A disk drive was replaced. 3) One or more disk drives were added. The problem is usually corrected automatically within 60 seconds, requiring no action from you. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

If the problem persists, check the cabling on or reseal any replaced or added disk drives, and check for failed disk drives. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for HA GROUP ERROR: DISK/SHELF COUNT MISMATCH

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.sfo.personality****Severity**

ALERT

**Description**

This message occurs when an All Flash Optimized node in a high-availability (HA) configuration detects its partner having a different personality setting. This mismatch might be due to either one node not being configured as All Flash Optimized, or the nodes being configured for different personalities (performance optimized or capacity optimized). Storage failover might be at risk of not functioning properly. HDDs taken over during failover events might not be able to serve data. If your system is configured to do so, it generates and transmits an AutoSupport® (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

An All-Flash Optimized node only recognizes SSDs and ignores HDDs. For the HA pair to be considered All-Flash Optimized, both nodes must be configured with the same personality (performance optimized or capacity optimized). Use the "system node show -instance" command to display each node's personality and optimization settings. If both nodes are configured with the same personality, all HDDs should be removed from the HA pair, because they will not be able to serve data. Contact NetApp technical support for assistance.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sfo.takeover

### Severity

ALERT

### Description

This message occurs when automatic takeover of the cluster partner is successful. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Address the issue that caused the takeover, then restore the cluster partner to normal operations. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sfo.takeover.failed

### Severity

ALERT

### Description

This message occurs when automatic takeover of the cluster partner is not successful. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Determine why the takeover did not succeed. Restore the cluster partner to normal operations. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for CONTROLLER TAKEOVER FAILED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sfo.takeover.halt

### Severity

NOTICE

### Description

This message occurs when takeover of the HA partner is successful. The takeover occurred because of an administrative halt on the HA partner. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and

resolution.

### **Corrective Action**

Address the issue that caused you to halt the HA partner and then restore it to normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CONTROLLER TAKEOVER COMPLETE HALT

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.sfo.takeover.m**

### **Severity**

NOTICE

### **Description**

This message occurs when manual takeover of the cluster partner is successful.

### **Corrective Action**

Address the issue that caused you to take over the cluster partner manually, then restore the cluster partner to normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CONTROLLER TAKEOVER COMPLETE MANUAL

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.sfo.takeover.m.dr**

### **Severity**

ERROR

### **Description**

This message occurs when manual takeover of the High Availability (HA) partner for disaster recovery (DR) is successful. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Address the issue that required you to perform a DR takeover of the HA partner. After the issue is resolved, restore the HA partner to normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CONTROLLER TAKEOVER COMPLETE MANUAL(DR)

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.sfo.takeover.panic**

### **Severity**

EMERGENCY

### **Description**

This message occurs when automatic takeover of the HA partner is successful. The controller takeover occurred because of a HA partner software failure or panic.

### **Corrective Action**

Address the cause of the cluster partner software failure or panic, then restore the cluster partner to normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CONTROLLER TAKEOVER COMPLETE PANIC

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.sfo.takeover.reboot**

### **Severity**

NOTICE

### **Description**

This message occurs when takeover of the HA partner is successful. The takeover occurred because of an administrative reboot on the HA partner. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Address the issue that caused you to reboot the HA partner and then restore it to normal operations. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for CONTROLLER TAKEOVER COMPLETE REBOOT

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.sfo.to.comm.loss**

### **Severity**

EMERGENCY

### **Description**

This message occurs when automatic takeover of the HA partner is successful. The controller takeover occurred because of a HA partner lost communication. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Address the cause of the cluster partner lost communication, then restore the cluster partner to normal operations. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.shlf events

### callhome.shlf.fan

#### Severity

EMERGENCY

#### Description

This message occurs when the system detects faulty hardware on the disk shelf, such as a fan, power supply unit (PSU), or failing temperature sensor. The problem might be environmental (temperature or faulty power) or hardware-related. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

#### Syslog Message

Call home for SHELF COOLING UNIT FAILED

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

### callhome.shlf.fan.warn

#### Severity

ERROR

#### Description

This message occurs when the system detects faulty hardware on the disk shelf, such as a fan, power supply unit (PSU), or failing temperature sensor. The problem might be environmental (temperature or faulty power) or hardware-related. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for SHELF COOLING UNIT FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.shlf.fault

### Severity

ERROR

### Description

This message occurs when the system detects faulty hardware on the disk shelf, such as a fan, power supply, or failing temperature sensor. The problem might be environmental (temperature or faulty power) or hardware-related. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C to 24C (65F to 75F). If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for SHELF\_FAULT

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.shlf.overtemp

### Severity

ERROR

### Description

This message occurs when the system detects a disk shelf problem. The problem might be due to environmental issues, such as temperature or faulty power, or by faulty hardware, such as a fan, power supply, or failing temperature sensor. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Evaluate the environment in which your system is operating and identify whether the problem is environmental or hardware-related. Your system should be in a room with an operating temperature of 18C

to 24C (65F to 75F). If faulty hardware caused the error, such as a bad temperature sensor or a broken fan, replace the faulty part as soon as possible. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.shlf.power.intr**

### **Severity**

ERROR

### **Description**

This message occurs when your system detects that a disk shelf power supply reported either a sustained loss of DC voltage output for about 20 seconds or a failure condition. This issue can occur for a number of reasons: 1. The power supply was turned off or unplugged. 2. The power to the server room failed or fluctuated. 3. The disk shelf power supply malfunctioned and should be replaced. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

If necessary, replace the disk shelf power supply. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for SHELF POWER INTERRUPTED

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.shlf.ps.fault**

### **Severity**

ERROR

### **Description**

This message occurs when your system detects that a disk shelf power supply reported a loss of DC voltage. This might be a transient issue. This issue can occur for a number of reasons: 1. The power supply was turned off or unplugged. 2. The power to the server room failed or fluctuated. 3. The shelf power supply malfunctioned and should be replaced. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

If necessary, replace the shelf power supply. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for SHELF POWER SUPPLY WARNING

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.shutdown events

## callhome.shutdown.pending

### Severity

ALERT

### Description

This message occurs when a disk drive fails but there are no suitable spares available for reconstruction. To protect your data, the system enters degraded mode. The system halts automatically to prevent a double disk drive failure, and possible loss of data, if it runs in degraded mode for the time interval given by the following command: "storage raid-options show -name raid.timeout -node local" The default timeout is usually 24 hours. If a spare drive becomes available while the system is running in degraded mode, the system immediately begins rebuilding the failed drive. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To prevent the system from shutting down in 24 hours, perform either of the following actions: 1. Add a spare disk drive suitable for reconstructing the failed disk drive. 2. Change the timeout period with the following command: "storage raid-options modify -name raid.timeout -node local -value X" where X equals the number of hours for the new timeout value. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SHUTDOWN PENDING (degraded mode)

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.sm events

## callhome.sm.vsd.break.fail

### Severity

ERROR

### Description

This message occurs when a break operation on a SnapMirror® relationship that is a part of VserverDR fails. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Contact NetApp technical support for additional information.



## Syslog Message

Call home for SNAPMIRROR BREAK FAILURE.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**SourceVserver** (STRING): Source Vserver of the SnapMirror relationship.

**DestVserver** (STRING): Destination Vserver of the SnapMirror relationship.

# callhome.snapmirror events

## callhome.snapmirror.oos

### Severity

NOTICE

### Description

This message occurs when a sync-granular consistency group (CG) transitions from in-sync status to out-of-sync status. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Look for the sm.syncmirror.out.of.sync event concerning this relationship in the EMS log file. The event provides the reason why the relationship has been transitioned to out-of-sync and an appropriate corrective action, if applicable.

## Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.snapmirror.oos.pers

### Severity

ALERT

### Description

This message occurs when a sync-granular consistency group (CG) relationship transitions from in-sync status to out-of-sync status and remains out-of-sync for longer than expected. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Look for the sm.syncmirror.out.of.sync event concerning this relationship in the EMS log file on source node of SnapMirror relationship. The event provides the reason why the relationship has been transitioned to out-of-sync. Verify network connections between the primary cluster and the secondary cluster. Check for error information in the "snapmirror show" command output for the relationship. Quiesce the relationship if the condition cannot be addressed immediately.

## Syslog Message

Call home for %s.

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.sp events

## callhome.sp.hbt.missed

### Severity

NOTICE

### Description

This message occurs when ONTAP® does not receive an expected "heartbeat" signal from the Service Processor (SP). Along with this message, log files from SP will be sent out for debugging. ONTAP will reset the SP to attempt to restore communication. The SP will be unavailable for up to two minutes while it reboots.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Call home for SP HBT MISSED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sp.hbt.stopped

### Severity

ALERT

### Description

This message occurs when ONTAP® is no longer receiving heartbeats from the Service Processor (SP). Depending on the hardware design, the system may continue to serve data or may determine to shut down to prevent data loss or hardware damage. The system continues to serve data, but because the SP might not be working, the system cannot send notifications of down appliances, boot errors, or Open Firmware (OFW) Power-On Self-Test (POST) errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If the system has shut down, attempt a hard power cycle: Pull the controller out from the chassis, push it back in then power on the system. Contact NetApp technical support if the problem persists after the power cycle, or for any other condition that may warrant attention.

## Syslog Message

Call home for SP HBT STOPPED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.spares events

## callhome.spares.low

### Severity

ERROR

### Description

This message occurs when one or more RAID groups do not have an eligible spare disk drive available in case of a disk drive failure. An eligible spare disk drive must be the correct size and capacity, and must be in the appropriate pool. If a disk drive fails without a spare to reconstruct on, the system enters degraded mode. The system halts automatically to prevent a double disk drive failure, and possible loss of data, if it runs in degraded mode for the time interval given by the following command: "storage raid-options show -name raid.timeout -node local" The default timeout is usually 24 hours. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Add suitable spare disk drives for all RAID groups to ensure that you have at least one spare disk drive. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for %s: %s

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**rg** (STRING): RAID groups on which spares low is identified.

# callhome.spm events

## callhome.spm.process.maxexit

### Severity

ERROR

### Description

This message occurs when a subsystem is not fully operational after one or more restarts and a restart threshold has been reached by the Service Process Manager (SPM). The inconsistency might be caused by external events such network data, previous configuration changes, software changes, firmware changes, or hardware changes. The message should not occur more than once per day. The underlying subsystem might be related to one or more operational systems such as management, authentication, network access, and so on. Performance issues and operational difficulties might have been exhibited before this event. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

If necessary, undo recent changes to the system and to external services that the system uses to stabilize operations, and then contact NetApp technical support. They will help to identify the underlying issue and whether a software upgrade has been made available to resolve the issue. You might need to isolate the system from any suspect network traffic.

### Syslog Message

Call home for MAXIMUM EXITS THRESHOLD OF A PROCESS MANAGED BY SPM IS REACHED

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.stackmon events

### callhome.stackmon.recon.failed

#### Severity

ALERT

#### Description

This message occurs when the system fails to assign the correct stack-id to one or more disk shelves. The shelf and disk names might not correctly reflect the storage topology. Shelf and disk operations might fail if multiple disks share the same name. If your system is configured to do so, it generates and transmits an AutoSupport(tm) (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### Corrective Action

Verify that the storage topology is correct. The (privilege: advanced) "storage stackmon repair" command can be used to diagnose and repair the issue. If you need further assistance, contact NetApp technical support.

#### Syslog Message

Call home for %s

#### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.svm events

### callhome.svm.mig.abort.fail

#### Severity

ERROR

#### Description

This message occurs when a Vserver migration failed to be aborted. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Run the "vserver migrate show" command to view the description of the error and the corrective action.

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## **callhome.svm.mig.aborted**

### **Severity**

NOTICE

### **Description**

This message occurs when a Vserver migration is successfully aborted. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

(None).

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## **callhome.svm.mig.clean.fail**

### **Severity**

ERROR

### **Description**

This message occurs when the cleanup phase of a Vserver migration has failed. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Run the "vserver migrate show" command to view the description of the error and the corrective action.

### **Syslog Message**

Call home for %s

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svm.mig.complete

### Severity

NOTICE

### Description

This message occurs when a Vserver migration has completed. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svm.mig.fail

### Severity

ERROR

### Description

This message occurs when a Vserver migration fails. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run the "vserver migrate show" command to view the description of the error and the corrective action.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svm.mig.pause.fail

### Severity

ERROR

### Description

This message occurs when a Vserver migration failed to be paused. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Run the "vserver migrate show" command to view the description of the error and the corrective action.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

**callhome.svm.mig.paused****Severity**

NOTICE

**Description**

This message occurs when a Vserver migration is paused. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

**callhome.svm.mig.pstCut.fail****Severity**

ERROR

**Description**

This message occurs when the post-cutover phase of a Vserver migration has failed. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Run the "vserver migrate show" command to view the description of the error and the corrective action.

**Syslog Message**

Call home for %s

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svm.mig.setCfg.fail

### Severity

ERROR

### Description

This message occurs when the setup configuration phase of a Vserver migration has failed. If the system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Run the "vserver migrate show" command to view the description of the error and the corrective action.

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svm.mig.started

### Severity

NOTICE

### Description

This message occurs when a Vserver migration has started successfully. If the system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for %s

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.  
**epoch** (STRING): Identifying tag for multi-node AutoSupport.

## callhome.svo events

### callhome.svo.prevented

### Severity

ERROR



## Description

This message occurs when a write to a SnapValidator enabled volume failed to be validated and was rejected by the server. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Check the configuration of the server to insure correct 'svf\_offset' setting on the lun (if applicable) and correct use of svo\_options. This includes making sure that RMAN data is not being stored on a volume that does not have the svo\_allow\_rman option set. If the server appears to be correctly configured, the problem may exist at any point between Oracle client and the server. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for SNAPVALIDATOR ERROR PREVENTED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.sync events

## callhome.sync.mirror.nospc

### Severity

ALERT

### Description

This message occurs when ONTAP® cannot create a new base Snapshot copy in a SyncMirror aggregate after multiple attempts. Systems using RAID SyncMirror require aggregate Snapshot copies for normal operation. Periodically, ONTAP creates a new base Snapshot copy with a name similar to "mirror\_resync.1201023128(filer\_name)." Once created, ONTAP deletes the old base Snapshot copy. The base Snapshot copy is used to optimize resynchronization when one plex goes offline. Short-lived circumstances when a new aggregate Snapshot copy cannot be created are not unusual. ONTAP initiates background processes to reclaim space and allow creation of new aggregate Snapshot copies. This message is issued when ONTAP cannot create a new base Snapshot copy after waiting a reasonable time. If the aggregate lacks an existing SyncMirror base Snapshot copy, then if one plex goes offline, subsequent resynchronization must copy all data in the aggregate - this can be a very long process. ONTAP issues this message when there is no SyncMirror base Snapshot copy and a new one cannot be created in repeated attempts in the last ten minutes (by default). If an old base Snapshot copy is present, this condition is less urgent. However, since space reclamation is likely to be slower, ONTAP waits up to two hours (by default) before issuing this message. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

ONTAP continually retries to create a new SyncMirror Snapshot copy. This condition might autocorrect, even after issuing this message. If it does not, contact NetApp technical support.

### Syslog Message

Call home for SYNCMIRROR SNAPSHOTS NOT ROTATING PROPERLY

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.syncm events

## callhome.syncm.isolation

### Severity

ALERT

### Description

This message occurs when both plexes of a SyncMirror® configuration have disk drives in the same hardware-based disk pool; that is, the plexes are not hardware-fault isolated. Therefore, a single hardware failure can cause both plexes to fail. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To resolve this issue, complete the following steps: 1. Identify the problem aggregate from the system logs. 2. Determine which mirror plexes have disk drives in mixed pools. 3. Determine how this occurred: possible causes include a wiring problem, a reconstruction forced mixed pools, or the mirror was created explicitly so that the plexes have disk drives in mixed pools. 4. Based on the information you gather, determine how to correct the issue. For example, recable the disk drives to restore hardware-fault isolation. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYNCMIRROR FAULT ISOLATION

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.syncm.plex

### Severity

ALERT

### Description

This message occurs when a SyncMirror® plex fails. Therefore, the SyncMirror plex is in a degraded mode. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem aggregate and failed plex from the system logs. 2. Determine the reason for the plex failure, such as disk drive, disk shelf, or loop failure. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYNCMIRROR PLEX FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.syncm.plex.offlined

### Severity

ALERT

### Description

This message occurs when a SyncMirror® plex is offlined. Therefore, the Syncmirror aggregate is in a degraded mode. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem aggregate and offlined plex from the system logs. 2. Determine the reason for the plex being offlined: a. Issue with backend array or adapter b. IO timeout errors on multiple disks c. SCSI errors 3. Replace any faulty components that you discover. 4. Online the plex manually. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYNCMIRROR PLEX OFFLINED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.syncm.verify

### Severity

ERROR

### Description

This message occurs when the system finds one or more block inconsistencies between plexes of a SyncMirror® configuration. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

To resolve this issue, complete the following steps: 1. Determine the problem aggregate and failed plex from the system logs. 2. Take the aggregate offline. 3. Split the plexes into two aggregates. 4. Contact NetApp technical support for assistance, and then run WAFL® Iron (wafliron) on both aggregates. 5. Determine which is the 'good' aggregate by inspection. 6. Take offline and then destroy the 'bad' aggregate. 7. Re-mirror the 'good' aggregate. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYNCMIRROR VERIFY FAILED

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.syncsm events

## callhome.syncsm.exception

### Severity

ERROR

### Description

This message occurs when critical thresholds are breached for a SnapMirror® Synchronous relationship. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Use the "snapmirror show" command to check for current status. Contact support personnel if this problem occurs repeatedly and prevents the relationship from moving to the InSync status.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.syncsm.oos.fatal

### Severity

ALERT

### Description

This message occurs when a resynchronize operation between the source volume and destination volume fails for a SnapMirror® Sync relationship. If your system is configured to do so, it generates and transmits an AutoSupport (or "call home") message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Look for the sms.resync.attempt.failed event concerning this relationship in the EMS log file. The event provides the reason why the relationship has been transitioned to out-of-sync and an appropriate corrective action, if applicable.

### Syslog Message

Call home for %s.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.sys events

## callhome.sys.config

### Severity

ERROR

### Description

This message occurs when the appliance has an unsupported PCI slot configuration. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Examine the command-line interface and System Configuration Guide for more information about slot assignments and hardware configuration, and then correct the unsupported configuration. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYSTEM CONFIGURATION WARNING

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.sys.config.crit

### Severity

ERROR

### Description

This message occurs when the appliance has an unsupported PCI slot configuration. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Examine the command-line interface and System Configuration Guide for more information about slot assignments and hardware configuration, and then correct the unsupported configuration. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for SYSTEM CONFIGURATION ERROR

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.tgt events

### callhome.tgt.lun.nospc

### Severity

ERROR

## Description

This message occurs when the SCSI Target subsystem begins failing SCSI Write commands due to an out of space (ENOSPC) condition on the appliance. This event is seen only when thin provisioning is used together with FCP or iSCSI. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

When this event occurs the customer should verify that all volumes containing luns on the appliance are properly configured for thin provisioning. Under some conditions additional disks may need to be added to the volume/aggregate. Contact NetApp support for more information. If you need assistance, contact NetApp technical support.

## Syslog Message

Call home for LUN OUT OF SPACE

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.unsupported events

## callhome.unsupported.snaplock

### Severity

ALERT

### Description

This message occurs when the system generates an AutoSupport event to notify you that a SnapLock® disk is attached to a system running a release of ONTAP® that does not support SnapLock.

### Corrective Action

Remove the SnapLock disks from this system and attach them to a system that supports SnapLock. You can also revert or upgrade this system to a release that supports SnapLock.

### Syslog Message

Call home for SNAPLOCK DISKS IN UNSUPPORTED RELEASE

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.ups events

## callhome.ups.line

### Severity

NOTICE

### Description

This message occurs when line power for an Uninterruptible Power Supply (UPS) system is restored and

the battery power is normal.

### **Corrective Action**

Maintain the environment and ensure that the line power to the UPS continues to be normal. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for UPS LINE POWER RESTORED

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ups.low.battery**

### **Severity**

ALERT

### **Description**

This message occurs when battery power for an Uninterruptible Power Supply (UPS) system has under 5 minutes of charge left. The system is currently serving data but will eventually shut down if the UPS continues to run without line power. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Check the environment and restore the line power to the UPS as soon as possible to avoid system shutdown. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for UPS LOW BATTERY

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.ups.ok.low.battery**

### **Severity**

NOTICE

### **Description**

This message occurs when line power for an Uninterruptible Power Supply (UPS) system is restored but the battery power is low. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Maintain the environment and ensure that the line power to the UPS continues to be normal. As line power continues to be normal, the battery is recharged. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for UPS PWR GOOD LOW BATTERY

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ups.ok.no.battery****Severity**

NOTICE

**Description**

This message occurs when line power for an Uninterruptible Power Supply (UPS) system is restored but the battery power is very low. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Maintain the environment and ensure that the line power to the UPS continues to be normal. As line power continues to be normal, the battery is recharged. If line power is disrupted or leaves the normal range, the system might shut down immediately to protect your data. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for UPS PWR GOOD CRITICAL BATTERY

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

**callhome.ups.on.battery****Severity**

ALERT

**Description**

This message occurs when line power for an Uninterruptible Power Supply (UPS) system fails and the system is running on battery power. The server is currently serving data but will eventually shut down if the UPS continues to run without line power. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

Check the environment and restore line power to the UPS as soon as possible to avoid system shutdown. If you need assistance, contact NetApp technical support.

**Syslog Message**

Call home for UPS ON BATTERY



## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.ups.shutting.down

### Severity

ALERT

### Description

This message occurs when battery power for an Uninterruptible Power Supply (UPS) system is almost depleted and the system is shutting down to protect your data. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Restore the line power to the UPS immediately to prevent system shutdown. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for UPS SHUTTING DOWN

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.user.events

### callhome.user.snmp.trap

### Severity

NOTICE

### Description

This message occurs when a SNMP (Simple Network Management Protocol) has been configured to cause an SNMP trap. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for SNMP USER DEFINED TRAP

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.vdom.events

## callhome.vdom.reconcile

### Severity

ERROR

### Description

This message occurs when the virtual disk infrastructure (vdom) unsuccessfully reconciles the Out Of Volume Container (OOVC) and the Vdisk Table Of Contents (VTOC). The Out Of Volume Container or OOVC is an area of storage that resides in a per-vvol aggregate metadir and provides storage for LUN configuration (override) and metadata (PR, mode page, log page). OOVC will have a LUN inventory that is identical to the LUN inventory in the AFS. Therefore the VTOC and OOVC must contain the same set of LUNs. As a result of data management operations like volume snapstore, the LUN inventory can differ between VTOC and OOVC. For these operations that affect LUN inventory, VDOM will need to reconcile the OOVC with the VTOC as part of the operation. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

(None).

### Syslog Message

Call home for %s Could not reconcile vtoc and oovc for volume associated with DSID: %llu, due to VDOM error: %u.

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**vdom\_error** (INT): The vdom error code associated with the reconcile.

## callhome.vfiler events

### callhome.vfiler.move.to.dst

### Severity

ALERT

### Description

This message occurs on the destination controller during the final phase of migrating a vfiler when the 120 second time limit has been reached . During the final phase of migrating a vfiler from a source controller to a destination controller, input/output data access is stalled until the destination is ready. The amount of time spent with I/O data access stalled within 120 seconds, since many pre-checks have been done to ensure this can be achieved. With the time limit reached, operations will resume to using the vfiler on the source controller. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check that service has been restored to the vfiler on the source controller. If not, contact NetApp technical support. Many pre-checks have been done to ensure timely migration of the vfiler from one controller to another. It is unclear whether network or other conditions have changed, such as changes in controller loads have prevented successful vfiler migration. Check the network connectivity, network bandwidth, load on the source controller and load on the destination controller. If they are not within normal expectations,

resolve the issue. If they are within normal expectations, consult with NetApp technical support. If you need further assistance, contact NetApp technical support.

### **Syslog Message**

Call home for VFILER MIGRATION EXCEEDED 120 SECONDS ON DESTINATION

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.vfiler.move.to.src**

### **Severity**

NOTICE

### **Description**

This message occurs on the source controller at the request of the destination controller. The destination controller has exceeded a 120 second time limit during the final phase of migrating a vfiler. During the final phase of migrating a vfiler from a source controller to a destination controller, input/output data access is stalled until the destination is ready. The amount of time spent with I/O data access stalled within 120 seconds, since many pre-checks have been done to ensure this can be achieved. With the time limit reached, operations will resume to using the vfiler on the source controller. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Check that service has been restored to the vfiler on this controller. If not, contact NetApp technical support. Many pre-checks have been done to ensure timely migration of the vfiler from one controller to another. It is unclear whether network or other conditions have changed, such as changes in controller loads have prevented successful vfiler migration. Check the network connectivity, network bandwidth, load on the source controller and load on the destination controller. If they are not within normal expectations, please resolve. If they are within normal expectations, consult with NetApp technical support. If you need further assistance, contact NetApp technical support.

### **Syslog Message**

Call home for VFILER MIGRATION TIME OUT EXCEEDED BY DESTINATION

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.vol events**

### **callhome.vol.offline**

### **Severity**

ALERT

### **Description**

This message occurs when the system detects an inconsistent volume and takes it offline. The file system inconsistency could be from one or more sources. The inconsistency might be related to previous operations or to current operations. Disabling access is a protective measure to ensure that further inconsistencies are avoided. If your system is configured to do so, it generates and transmits an

AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

If the operating environment for the system, storage, or associated cabling has changed in terms of temperature, humidity, vibration, unstable electrical power, or other physical factors, gracefully shut down and power off the system until the environment is restored to normal operations. If the operating environment has not changed, identify the source of the inconsistency by checking for previous errors and warnings. Also check for hardware statistics from Fibre Channel, SCSI, disk drives, other communications mechanisms, and previous administrative activities. Another approach is to consult with NetApp technical support (it is recommended that you do so before attempting this step), and then run WAFL® Iron (wafliron) on the volume to identify and optionally repair the inconsistency. If you need further assistance, contact NetApp technical support.

### **Syslog Message**

Call home for VOLUME INCONSISTENT AND TAKEN OFFLINE

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.vol.space.crit**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a volume has overallocated space. It will be taken offline soon and the node will reboot. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Create space by increasing volume/aggregate size or deleting data. While changing the provisioning from thick to thin will prevent further utilization of space, it is no longer sufficient to prevent the volume/aggregate from being offlined. In a volume, size can be increased with the "volume size" command. To change provisioning in a volume, reserved files can be unreserved using the "volume file reservation" command. In an aggregate, size can be increased by adding disks. To change provisioning in an aggregate, volume guarantees can be changed from "volume" to "none" using the space-guarantee parameter in the "volume modify" command. The nodeshell "vol status -S" and "aggr status -S" commands display "Total Used" and "Total Physical Used" values greater than 100% when the reserve is being used. The "vol.log.overalloc" and "volume.phys.overalloc" EMS events and their associated SNMP traps of "volumeLogicalOverallocated" and "volumePhysicalOverallocated" are issued when the volume or aggregate is consuming more than the recommended amount of reserve.

### **Syslog Message**

Call home for VOLUME CRITICALLY OUT OF SPACE ERROR.

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

# callhome.wafl events

## callhome.wafl.check

### Severity

ALERT

### Description

This message occurs after a WAFL® check (wack) operation runs. WAFL check is a detection and repair process. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Check the appliance's log for valuable information that WAFL check might have reported. If you need assistance, contact NetApp technical support.

### Syslog Message

Call home for WAFL\_check

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.wafl.inconsistent

### Severity

ALERT

### Description

This message occurs when the system detects that a volume is inconsistent. The file system inconsistency can be from one or more sources. The inconsistency can be related to previous operations or to current operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### Corrective Action

Consult with NetApp technical support (it is recommended that you do so before attempting this step), and then run WAFL® Iron (wafliron) on the volume to identify and optionally repair the inconsistency. If you need further assistance, contact NetApp technical support.

### Syslog Message

Call home for WAFL INCONSISTENT

### Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.wafl.inconsistent.block

### Severity

ALERT

## Description

This message occurs when the system detects a volume that contains an inconsistent block. The file system inconsistency can be from one or more sources. The inconsistency can be related to previous data operations or to current data operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Contact NetApp technical support for further assistance.

## Syslog Message

Call home for WAFL INCONSISTENT BLOCK

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.wafl.inconsistent.user.block

## Severity

ALERT

## Description

This message occurs when the system detects inconsistent user data within a volume. The file system inconsistency can be from one or more sources. The inconsistency can be related to previous data operations or to current data operations. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

## Corrective Action

Contact NetApp technical support for further assistance.

## Syslog Message

Call home for WAFL INCONSISTENT USER BLOCK

## Parameters

**subject** (STRING): AutoSupport subject or title for this event.

## callhome.wafl.iron.done

## Severity

NOTICE

## Description

This message occurs after a WAFL® Iron (wafiliron) operation runs. WAFL® Iron is a detection and repair process. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

### **Corrective Action**

Check the appliance's log for valuable information that WAFL Iron might have reported. If you need assistance, contact NetApp technical support.

### **Syslog Message**

Call home for WAFL IRON COMPLETED

### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.wafliron events**

### **callhome.wafliron.chg.pending**

#### **Severity**

NOTICE

#### **Description**

This message occurs when WAFL® Iron (wafliron) Optional Commit changes are pending. WAFL® Iron is a detection and repair process. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

#### **Corrective Action**

Work with NetApp technical support on this issue. If you decide against this, it is recommended that the pending changes be committed.

#### **Syslog Message**

Call home for WAFLIRON CHANGES PENDING

#### **Parameters**

**subject** (STRING): AutoSupport subject or title for this event.

## **callhome.weekly events**

### **callhome.weekly**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the weekly AutoSupport or 'call home' event is sent. Information is sent to NetApp technical support to ensure that subsequent communications work and to enable operations analysis. Some results of these analysis might include recommendations on upgrades, firmware changes, and configuration changes to improve the reliability and performance of the system. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.

**Corrective Action**

(None).

**Syslog Message**

Call home for WEEKLY\_LOG

**Parameters**

**subject** (STRING): AutoSupport subject or title for this event.



# cam events

## cam.dev events

### cam.dev.fuaNotSupported

#### Severity

ALERT

#### Description

This message occurs when the disk does not support the Forced Unit Access (FUA) feature bit. Using this disk with the write cache enabled can lead to data corruption.

#### Corrective Action

Use disks with the write cache disabled or disks where the FUA feature is supported.

#### Syslog Message

FUA feature is not supported for disk %s%d with target address %s.

#### Parameters

**periph\_name** (STRING): Disk name.

**unit** (INT): Unit number.

**target\_addr** (STRING): Disk target address.

### cam.dev.fuaNotVerified

#### Severity

ALERT

#### Description

This message occurs when ONTAP® software fails to read the SCSI page containing the Forced Unit Access (FUA) feature bit. Using the disk with the write cache enabled and FUA not supported can lead to data corruption.

#### Corrective Action

Because ONTAP cannot determine the state of the FUA feature bit, disable the disk write cache if it is enabled by default. If the write cache cannot be disabled, replace the disk to avoid corruption.

#### Syslog Message

Unable to read the FUA settings for disk %s%d with target address %s.

#### Parameters

**periph\_name** (STRING): Disk name.

**unit** (INT): Unit number.

**target\_addr** (STRING): Disk target address.

## cam.timeout events

## **cam.timeout.retry**

### **Severity**

NOTICE

### **Description**

This message occurs when an I/O request to a Common Access Method (CAM) device driver times out. This might be indicative of an overloaded back-end storage subsystem. The I/O operation will be retried; this message does not indicate failure.

### **Corrective Action**

(None).

### **Syslog Message**

CAM device driver I/O timeout. %s%s%s.

### **Parameters**

**cam\_timeout** (STRING): Timeout message for an I/O request from a CAM device driver.

**uuid** (STRING): Unique identifier for the device.

**outstanding** (STRING): Command that timed out and its elapsed time.

# cap events

## cap.cache events

### cap.cache.corrupt

#### Severity

ERROR

#### Description

This message occurs when a corrupt Central Access Policy (CAP) is encountered during access checks.

#### Corrective Action

Fix the CAP definitions on the Active Directory, and use the 'vserver cifs group-policy update' command to refresh the cache on the cluster.

#### Syslog Message

Encountered a corrupt CAP for Vserver %s.

#### Parameters

**vserverUuidStr** (STRING): UUID of the Vserver.

### cap.cache.not.ready

#### Severity

ERROR

#### Description

This message occurs when the system attempts to access a Central Access Policy (CAP) for evaluation during access checks and finds the CAP cache database (DB) is not ready for access.

#### Corrective Action

Make sure that group-policy is enabled for the given Vserver and the Central Access Policies (CAPs) from the Active Directory (AD) can be applied to file system objects.

#### Syslog Message

CAP Cache DB not ready for Vserver %s.

#### Parameters

**vserverUuidStr** (STRING): UUID of the Vserver.

### cap.cache.update.failed

#### Severity

ERROR

#### Description

This message occurs when refresh of the Central Access Policy (CAP) Cache database (DB) fails.

**Corrective Action**

Use the "vserver cifs group-policy update" command with the "force" flag to retry the operation manually.

**Syslog Message**

Failed to update the CAP cache DB for Vserver %s. (Reason: %s.)

**Parameters**

**vserverUuidStr** (STRING): UUID of the Vserver.

**errString** (STRING): Error string.

# ccma events

## ccma.disabled events

### ccma.disabled

#### Severity

ERROR

#### Description

This message occurs when the Performance Archive has been disabled for more than an hour. This is caused by setting the `-is-enabled` flag to false. This message will be repeated every hour that the archiver is turned off. The messages will only displayed once a day.

#### Corrective Action

Run the `'statistics archive config modify -is-enable true'` command. This will enable the Performance Archive.

#### Syslog Message

The Performance Archive has been disabled for more than an hour. It is Important for the analysis and diagnoses of performance related issues. Consider enabling it.

#### Parameters

(None).

## ccma.perfstatd events

### ccma.perfstatd.disabled

#### Severity

ERROR

#### Description

This message occurs when the Perfstat daemon (`perfstatd`) has been disabled for more than an hour. This is caused by setting the `-perfstat-sampling-period` value to 0. This message will be repeated every hour `perfstatd` is turned off. The messages will only displayed once a day.

#### Corrective Action

Run the `'statistics archive config modify -perfstat-sampling-period 5'` command. This will enable the `perfstatd` and it will set the Perfstat sampling period to the default value of 5 minutes.

#### Syslog Message

The Perfstat daemon (`perfstatd`) has been disabled for more than an hour. It is important for the analysis and diagnoses of performance related issues. Consider enabling it.

#### Parameters

(None).

# ccma.quota.events

## ccma.quota.datastore.leak

### Severity

ERROR

### Description

This message occurs when a datastore has been orphaned due to the system not having sufficient memory to keep track of it in the quota system. Orphaned datastores will take up disk space and may build up over time if this problem recurs.

### Corrective Action

Delete the orphaned datastore and/or restart the CM daemon once the out of memory condition has been resolved.

### Syslog Message

A datastore may have been orphaned due to the system being out of memory. Please delete this datastore: %s

### Parameters

**datastore** (STRING): Name of the datastore that has been orphaned and needs to be deleted.

## ccma.quota.file.leak

### Severity

ERROR

### Description

This message occurs when an archiver file has been orphaned due to the system not having sufficient memory to keep track of the file in the quota system. Orphaned files will take up disk space and may build up over time if this problem recurs.

### Corrective Action

Delete the orphaned file and/or restart the CM daemon once the out of memory condition has been resolved.

### Syslog Message

An archiver file may have been orphaned due to the system being out of memory. Please delete this file: %s

### Parameters

**filename** (STRING): Name of the archiver file that has been orphaned and needs to be deleted.

## ccma.quota.throughput

### Severity

NOTICE

### Description

This message occurs when the performance archive data collection rate is too high for the amount of reserved disk space, possibly resulting in a shortened history of the system's performance (less than the

configured `datastore-max-retention` value). A shortened history might limit the ability of support staff to triage performance problems.

### Corrective Action

Using the (diagnostic privilege) "statistics archive config" commands, increase the performance archive's `datastore-max-percent` configuration parameter so that the reported `datastore-max-size` value on this node is increased by the stated "size" amount.

### Syslog Message

The performance archive data collection rate is too high for the amount of space available. The system can only keep %ld days of performance archive data at this rate, and needs %ld MB of additional space to maintain a system performance history equal to the `datastore-max-retention` configuration value.

### Parameters

**days** (LONGINT): Number of days of performance archive data that can be stored at the current rate.

**size** (LONGINT): Size increase (in MB) needed to maintain a system performance history equal to the `datastore-max-retention` configuration value.

# cde events

## cde.undo events

### cde.undo.abort

**Severity**

NOTICE

**Description**

This message occurs when undo of cde has to abort.

**Corrective Action**

(None).

**Syslog Message**

Undo of auto adaptive compression on volume %s%s aborted.

**Parameters**

**vvol\_name** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

### cde.undo.done

**Severity**

NOTICE

**Description**

This message occurs when cde undo is done.

**Corrective Action**

(None).

**Syslog Message**

Auto adaptive compression undo is done for volume %s%s.

**Parameters**

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.



# cdpd events

## cdpd.recv events

### cdpd.recv.cacheOverflow

#### Severity

ERROR

#### Description

This message occurs when the CDP (Cisco Discovery Protocol) daemon running on the storage controller receives more unique neighbor advertisements than it can cache. The storage controller stops saving incoming CDP advertisements after the CDP neighbor cache has the maximum allowed neighbor entries. The CDP daemon will start processing new advertisements again after entries in the CDP neighbor cache start timing out.

#### Corrective Action

Use the 'cdpd show-stats' CLI command under the "Cache overflow" receive statistic to view the number of unique advertisements in excess of the maximum. Check the traffic statistics on the neighboring switches and bridges. If the storage controller has more than 64 CDP neighbors, you can ignore this message. Otherwise, take a packet trace on the network or on each network port of the storage controller to check the source of the CDP advertisements.

#### Syslog Message

The CDP neighbors table is full.

#### Parameters

(None).

# cecc\_log events

## cecc\_log.dropped events

### cecc\_log.dropped

#### Severity

NOTICE

#### Description

This message occurs when the ECC error logger dropped some correctable ECC errors without logging them. Typically, correctable ECC errors are logged in the system even though they are automatically corrected. They are dropped due to lack of buffer space in the ECC error logger. This might be a sign of high frequency of ECC errors. If the number of ECC errors dropped crosses a certain predefined threshold, the ECC error logger issues a `cecc_log.masked` message.

#### Corrective Action

(None).

#### Syslog Message

Dropped %lld additional correctable ECC error details.

#### Parameters

`num_cecc_batch_dropped` (LONGINT): Number of correctable ECC errors dropped without being logged.

## cecc\_log.entry events

### cecc\_log.entry

#### Severity

NOTICE

#### Description

Correctable ECC errors occurred in the system memory. The system automatically corrects these errors. All the corrected errors are reported by the ECC error logger to track the history of occurrence of these errors.

#### Corrective Action

(None).

#### Syslog Message

%lld: %s

#### Parameters

`cecc_entry_num` (LONGINT): Number of correctable ECC errors since boot.

`cecc_msg` (STRING): Information about the ECC error.

## cecc\_log.entry\_no\_syslog events

## **cecc\_log.entry\_no\_syslog**

### **Severity**

NOTICE

### **Description**

Correctable ECC errors occurred in the system memory. The system automatically corrects these errors. All the corrected errors are reported by the ECC error logger to track the history of occurrence of these errors.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**cecc\_entry\_num** (LONGINT): Number of correctable ECC errors since boot.

**cecc\_msg** (STRING): Information about the ECC error.

## **cecc\_log.masked events**

### **cecc\_log.masked**

### **Severity**

ERROR

### **Description**

This message occurs when the ECC error logger stops reporting the occurrence of correctable ECC errors because the frequency of their occurrence is too high. The ECC logger crossed the predefined threshold of ECC errors that can be dropped. Reporting is suppressed to control the log traffic, but the system sends out a periodic message as a reminder.

### **Corrective Action**

Use the "system controller memory dimm show" command to check system memory.

### **Syslog Message**

Error logging temporarily suspended due to excessive CECC errors.

### **Parameters**

(None).

## **cecc\_log.summary events**

### **cecc\_log.summary**

### **Severity**

NOTICE

### **Description**

This message occurs when the ECC error logger reports its periodic summary of new correctable ECC

errors that occurred after the last report and the total number of ECC errors that occurred since boot. These errors were corrected by the system.

### **Corrective Action**

(None).

### **Syslog Message**

Total of %lld new correctable ECC errors just reported. You might want to check system memory. %lld correctable ECC errors reported since booting.

### **Parameters**

**num\_ceccs** (LONGINT): Number of new correctable ECC errors after the last summary.

**total\_num\_ceccs** (LONGINT): Number of correctable ECC errors since boot.

# certificate events

## certificate.insecure events

### certificate.insecure.hash

#### Severity

ERROR

#### Description

This message occurs when an installed certificate uses a hash function that is no longer considered secure. Client-server communication using this certificate should be presumed insecure, and could be disrupted.

#### Corrective Action

Delete the insecure certificate by using the "security certificate delete" command. Install an appropriate replacement by using the "security certificate install" command or "security certificate create" command.

#### Syslog Message

A digital certificate with UUID %s in Vserver %s with name %s is using the insecure hash function %s.

#### Parameters

**uuid** (STRING): UUID of the offending certificate.

**vserver** (STRING): Name of the Vserver in which the offending certificate is installed.

**cert\_name** (STRING): Name (unique within a Vserver) of the offending certificate.

**hash\_function** (STRING): Hash function used by the offending certificate.

### certificate.insecure.key

#### Severity

ERROR

#### Description

This message occurs when an installed certificate uses a key length that is no longer considered secure for the corresponding key algorithm. Client-server communication using this certificate should be presumed insecure, and could be disrupted.

#### Corrective Action

Delete the insecure certificate by using the "security certificate delete" command. Install an appropriate replacement by using the "security certificate install" command or "security certificate create" command.

#### Syslog Message

A digital certificate with UUID %s in Vserver %s with name %s has an insecure key of type %s and length %d.

#### Parameters

**uuid** (STRING): UUID of the offending certificate.

**vserver** (STRING): Name of the Vserver in which the offending certificate is installed.

**cert\_name** (STRING): Name (unique within a Vserver) of the offending certificate.

**key\_algorithm** (STRING): Algorithm (e.g., RSA) of the key for the corresponding certificate.

**key\_length** (INT): Length of the key in bits for the corresponding certificate.

# certificate.invalid events

## certificate.invalid.chain

### Severity

ERROR

### Description

This message occurs when an installed certificate record contains multiple X.509 certificates. Such records are not treated consistently by the system, but were unintentionally permitted in prior releases.

### Corrective Action

Delete the problematic record by using the "security certificate delete" command. Install the constituent certificates separately by using the "security certificate install" command once for each certificate.

### Syslog Message

The certificate record with UUID %s, named "%s" in Vserver "%s", contains multiple certificates.

### Parameters

**uuid** (STRING): UUID of the offending record.

**cert\_name** (STRING): Name (unique within a Vserver) of the offending record.

**vserver** (STRING): Name of the Vserver in which the offending record is installed.

# certificate.node events

## certificate.node.vserver

### Severity

ERROR

### Description

This message occurs when an installed certificate record exists in a node Vserver. Such records are not treated consistently by the system, but were unintentionally permitted in prior releases.

### Corrective Action

Delete the problematic record by using the "security certificate delete" command.

### Syslog Message

The certificate record with UUID %s, named "%s" is installed in the node Vserver "%s".

### Parameters

**uuid** (STRING): UUID of the offending record.

**cert\_name** (STRING): Name (unique within a Vserver) of the offending record.

**vserver** (STRING): Name of the node Vserver in which the offending record is installed.

# cf events

## cf.autogbcanceled events

### cf.autoGBCanceled.headswap

**Severity**

INFORMATIONAL

**Description**

This message occurs when an automatic giveback is canceled because of a headswap.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

## cf.badshelfcountmsg events

### cf.badShelfCountMsg

**Severity**

ERROR

**Description**

This event occurs if we are sent an invalid negotiated failover shelf count message. The message size is incorrect. This is likely due to an InterConnect problem.

**Corrective Action**

Check for and resolve any interconnect errors. If the nodes are on different releases, upgrade the down-rev node.

**Syslog Message**

Invalid disk shelf count message was received from partner.

**Parameters**

(None).

## cf.badshelfcountversion events

### cf.badShelfCountVersion

**Severity**

ERROR

## Description

This event occurs if we are sent a negotiated failover shelf count message with a version which we don't support. As a result, disk shelf count negotiated failover will not work. This is likely due to a kernel version mismatch.

## Corrective Action

Upgrade the down-rev node so both nodes are on the same version.

## Syslog Message

Disk shelf count message with unsupported version received from partner.

## Parameters

(None).

# cf.boot events

## cf.boot.continuing

### Severity

INFORMATIONAL

### Description

This message occurs when the failover monitor continues booting after waiting for giveback.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

(None).

# cf.cdo events

## cf.cdo.badMsg

### Severity

ERROR

### Description

This event is issued when the filer receives a disk operation message over the High Availability (HA) interconnect that it does not recognize. This is likely due to a change in the message protocol that the down-revision partner does not recognize. Fix this problem by upgrading the partner to a compatible version of Data ONTAP®. This message is not emitted for every message received to avoid spamming the console.

### Corrective Action

Upgrade both nodes to the same release.



## Syslog Message

unknown message %d received from partner (%d unknown messages since boot).

## Parameters

**msg** (INT): Message which was not known.  
**count** (INT): Number of unknown messages.

## cf.cdo.reportOp

### Severity

NOTICE

### Description

This event is issued when the filer bypasses or unbypasses a disk after being requested to do so by its partner. The bypassing of a drive on a particular port may only be done by the host which is attached to that port; the High Availability (HA) software ensures that both ports of the drive are bypassed by communicating bypass requests to the partner.

### Corrective Action

(None).

## Syslog Message

%s disk %s on behalf of the partner.

## Parameters

**op** (STRING): Operation description  
**driveName** (STRING): Drive name

## cf.disk events

### cf.disk.invent.mismatchalt

### Severity

ALERT

### Description

This message occurs when one of the nodes in a high-availability (HA) pair reports that multiple disks in the disk inventory do not match what the partner node reported. This only occurs when sufficient mismatches have occurred to throttle the specific disk inventory mismatch message. The issue might be due to one of following reasons: (1) One node can see disks that the other node cannot. (2) Ownership of multiple disks has changed. (3) Multiple disks have been failed or unfailed. (4) Multiple disks have been inserted or removed.

### Corrective Action

Rescan the disks on both nodes in the HA pair to determine the latest ownership by using the "storage disk refresh-ownership" command. If the inventory mismatch is not resolved, remove and replace any failed disks. If there are no failed disks, check the cabling for loose connections or a bad cable on the node where the disks are not included in the inventory.

## Syslog Message

Status of some of the disks has changed or the node (%s) is missing %u disks (detailed logs have been throttled).

## Parameters

**host** (STRING): Name of the node not reporting the disks in its disk inventory.  
**nummissing** (INT): Number of disks not reported in the disk inventory.

## cf.disk.invent.mismatchOKalt

### Severity

INFORMATIONAL

### Description

This message occurs when previously detected disk inventory mismatches have been resolved and the detailed match events have been throttled.

### Corrective Action

(None).

## Syslog Message

%u disks now included in the inventory of node (%s) (detailed logs have been throttled).

## Parameters

**nummatches** (INT): Number of disks found to match the partner disk inventory so far.  
**host** (STRING): Name of the host where the disk is included in the inventory.

## cf.disk.inventory.mismatch

### Severity

ERROR

### Description

This message occurs when one of the nodes in a high-availability (HA) pair has reported this disk in its disk inventory, but the HA partner node has not. This might be due to one of following reasons: (1) One node can see the disk, but the other node cannot. (2) Ownership of the disk has changed. (3) The disk has either been failed or unfailed. (4) The disk has been inserted or removed.

### Corrective Action

Use the "storage disk refresh-ownership" command on both nodes in the HA pair to rescan the disks and to determine the latest ownership. If the inventory mismatch is not resolved, remove and replace any failed disks. If there are no failed disks, check the cabling for loose connections or a bad cable on the node where the disk is not included in the inventory.

## Syslog Message

Status of the disk %s (%s) has recently changed or the node (%s) is missing the disk.

## Parameters

**diskname** (STRING): Name of the disk not reported in the inventory.  
**uid** (STRING): Unique identifier of the disk that was not included in the inventory.  
**host** (STRING): Name of the node not reporting the disk in its disk inventory.

## **cf.disk.inventory.mismatchOK**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a previously detected disk inventory mismatch is resolved.

### **Corrective Action**

(None).

### **Syslog Message**

The node (%s) included the disk %s (%s) in its inventory.

### **Parameters**

**host** (STRING): Name of the host where the disk is included in the inventory.

**diskname** (STRING): Name of the disk included in the inventory.

**uid** (STRING): Unique identifier of the disk.

## **cf.disk.inventory.ver.notok**

### **Severity**

ERROR

### **Description**

This message occurs when the local node does not support the partner's disk inventory message version.

### **Corrective Action**

Upgrade the partner to the same Data ONTAP® version.

### **Syslog Message**

cfdisk: The local node did not support the partner's disk inventory message version. The disk inventory check will not be performed during giveback.

### **Parameters**

(None).

## **cf.disk.inventoryChange**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when High Availability (HA) disk inventory is enabled/disabled.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: disk inventory %s

## Parameters

**state** (STRING): Parameter which describes if enabled or disabled.

## cf.disk.releaseFailed

### Severity

ERROR

### Description

This message occurs when the release of reservation request on a disk fails. The error indicates that a disk is not ready, that it failed, or that it does not exist.

### Corrective Action

Use the "disk show -broken" CLI command to check whether the disk is reporting a "FAILED" state, or use the "disk show -disk diskname" command to check the status of the disk. Also check the syslog for other errors pertaining to the same disk. The disk might need to be reseated or replaced.

### Syslog Message

Disk reservation failed to release on disk %s because the error %d was encountered.

## Parameters

**diskname** (STRING): Disk that reported the failure when a release of reservation request was attempted on it.

**force** (INT): Type of reservation release request. If non-zero, then the system was instructed to release reservations owned by the other systems. Otherwise, it only releases the reservation on the disks reserved by the local system.

**error** (INT): Error encountered while attempting to release the reservation on disk.

**errortext** (STRING): Error description as text.

## cf.disk.reserveFailed

### Severity

ERROR

### Description

This message occurs when a disk reserve fails. The error indicates that a disk is not ready, that it failed, or that it does not exist.

### Corrective Action

Check the specified disk for failure. You might need to replace it or reseat it. Then check the syslog for other errors pertaining to the same disk.

### Syslog Message

(None).

## Parameters

**diskname** (STRING): Disk that reported the failure when a reserve was attempted on it.

**preempt** (INT): Type of reservation request. If non-zero, then the system was instructed to preempt reservations from the other systems. Otherwise, it only reserves disks without existing reservations.

**error** (INT): Error encountered while attempting to reserve the disk.

**errortext** (STRING): Error description as text.

## cf.disk.resumeDisks

### Severity

NOTICE

### Description

This message occurs when a node resumes partner disks after a takeover. If any disks fail to be resumed, the node makes another attempt to resume them after short delay.

### Corrective Action

(None).

### Syslog Message

Node attempted to resume %d disks, but failed with %d disks.

### Parameters

**fmdisk\_count** (INT): Number of disks resumed.

**fmdisk\_error\_count** (INT): Number of disks that had resume errors.

## cf.disk.ResvFail

### Severity

ALERT

### Description

This message occurs when a node attempts to access a disk that has been reserved by the HA partner node. A node performing a takeover reserves the HA partner node's disks as part of a takeover operation. It is possible that the node being taken over can attempt to access a reserved disk before it has shut down as part of being taken over. The takeover node will automatically release the reservation prior to a giveback operation.

### Corrective Action

The takeover node will automatically release the disk reservation before a giveback operation.

### Syslog Message

Disk %s has been reserved by the High Availability (HA) partner as part of a takeover operation.

### Parameters

**DiskName** (STRING): Name of the disk.

## cf.disk.ResvTakeOver

### Severity

NOTICE

### Description

This event is issued when we find disks with reservation conflicts.

### Corrective Action

(None).

### **Syslog Message**

This node will wait for giveback and the disk reservations to be released.

### **Parameters**

(None).

## **cf.disk.skipped**

### **Severity**

NOTICE

### **Description**

This message occurs when the failover monitor skips an unresponsive disk.

### **Corrective Action**

Use the "storage disk show" command to check whether the disk failed. Replace the disk if needed.

### **Syslog Message**

The disk %s was skipped because it reported the status %s.

### **Parameters**

**diskname** (STRING): Name of the disk that was skipped.

**status** (STRING): Status reported by the disk.

## **cf.diskinventory events**

### **cf.diskinventory.cksum.fail**

### **Severity**

ALERT

### **Description**

This message occurs when the computed checksum of the disk inventory message is different from the checksum sent as part of the disk inventory message.

### **Corrective Action**

(None).

### **Syslog Message**

Disk inventory checksum mismatch: old %d new %d.

### **Parameters**

**oldCkSum** (INT): Old checksum, sent as part of the message.

**newCkSum** (INT): Checksum of the disk inventory message computed locally.

### **cf.diskinventory.many.disks**

### **Severity**

INFORMATIONAL

## Description

This message occurs when a node receives more than maximum supported disks in the disk inventory message.

## Corrective Action

(None).

## Syslog Message

disk inventory: ndisks %d max\_disks %d

## Parameters

**ndisks** (INT): Number of disk descriptions sent in the inventory message.

**maxdisks** (INT): Maximum number of disks supported.

# cf.fm events

## cf.fm.cpuUtilDuringTOAndGB

### Deprecated

Deprecated as of version 9.0.

### Severity

NOTICE

### Description

This message occurs at the start of a takeover, end of a successful takeover, start of a CFO giveback, and completion of an SFO giveback. It records the maximum, minimum, and average CPU and disk utilization on the node executing the takeover or giveback.

### Corrective Action

(None).

### Syslog Message

CPU and disk utilization during the %d seconds %s: cpu\_util\_high: %lld; cpu\_util\_low: %lld; cpu\_util\_avg: %lld; disk\_util\_high: %lld; disk\_util\_low: %lld; disk\_util\_avg: %lld

### Parameters

**window\_sz** (INT): Duration, in seconds, over which CPU and disk utilization are tracked.

**when** (STRING): Event during which CPU and disk utilization are tracked.

**cpu\_util\_high** (LONGINT): Maximum CPU utilization.

**cpu\_util\_low** (LONGINT): Minimum CPU utilization.

**cpu\_util\_avg** (LONGINT): Average CPU utilization.

**disk\_util\_high** (LONGINT): Maximum disk utilization.

**disk\_util\_low** (LONGINT): Minimum disk utilization.

**disk\_util\_avg** (LONGINT): Average disk utilization.

## cf.fm.discardNvram

### Severity

NOTICE

**Description**

This event is issued when we discover that the partner has previously taken us over, forcing us to invalidate our own nvram contents. This is a normal condition, subsequent to a takeover/giveback.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: node was previously taken over, nvram may be discarded

**Parameters**

(None).

**cf.fm.diskInventoryOff****Severity**

ERROR

**Description**

This message occurs when the system discovers that disk inventory gathering has been disabled. During normal operation, the high-availability (HA) nodes transmit their disk inventory data at regular intervals. This is intended to prevent a situation in which loop connectivity problems are unnoticed until a takeover event occurs. If this event occurs, contact NetApp technical support.

**Corrective Action**

Use the "sysconfig" and "storage" nodeshell commands to determine whether there are problems with the loop, adapter, or shelf. Resolve those problems.

**Syslog Message**

Failover monitor: HA disk inventory disabled.

**Parameters**

(None).

**cf.fm.diskRelease****Severity**

INFORMATIONAL

**Description**

This event is issued when we're using a debug build and failover monitor reservations are released.

**Corrective Action**

n/a

**Syslog Message**

Failover monitor: released disk reservations.

**Parameters**

(None).



## **cf.fm.diskReleaseFail**

### **Severity**

NOTICE

### **Description**

This message occurs when the release of a reservation on a disk fails in preparation for a giveback event. The error indicates that a disk is not ready, that it failed, or that it does not exist. If the reservation is detected by the partner node, it will reboot.

### **Corrective Action**

Look for the cf.disk.releaseFailed event in the EMS log to find the name of the disk where the reservation could not be released. Follow the corrective action described in the cf.disk.releaseFailed event to address any problems with the disk.

### **Syslog Message**

Could not release disk reservations of at least one disk.

### **Parameters**

(None).

## **cf.fm.duplicateId**

### **Severity**

ALERT

### **Description**

This message occurs when the local node system identifier is the same as the partner's. This could happen if the HA-Interconnect is configured for loopback in maintenance mode systems or if the system was not properly configured. The local node will halt in this case and the partner node will do a takeover of the local node resources, provided takeover is enabled.

### **Corrective Action**

If this message occurs only while the system is configured in maintenance mode, it can be ignored as HA-interconnect loopback tests send a message with a node's own system identifier to itself. If this message occurs while a system is not in maintenance mode, check if the HA-interconnect cables are properly connected. If cabling is correct, contact NetApp technical support for assistance.

### **Syslog Message**

Partner ID %u is the same as that of this node. This node will halt and the partner will perform a takeover, if takeover is enabled.

### **Parameters**

id (INT): System ID.

## **cf.fm.earlyGivebackDone**

### **Severity**

NOTICE

## Description

This event occurs when we are aborting a takeover that was initiated during a previous boot sequence. This event should only occur under unusual circumstances, indicating successful recovery from a software failure.

## Corrective Action

(None).

## Syslog Message

Failover monitor: giveback of previous takeover complete

## Parameters

(None).

## cf.fm.earlyTakeoverFailed

### Severity

ALERT

### Description

This message occurs when an error during early takeover prevents the node from booting into takeover mode. The node instead boots up without taking over its partner and also releases its partner resources, allowing the partner node to boot up. Note: Early takeover occurs when a node boots up after rebooting while in takeover mode.

### Corrective Action

Check the EMS log for the cf.rsrc.takeoverFail error or other errors indicating why the node could not boot into takeover mode.

### Syslog Message

Early takeover failed; node will boot without taking over partner node. Partner resources released, allowing partner node to boot.

### Parameters

(None).

## cf.fm.fastTimeoutBlocked

### Severity

ERROR

### Description

This event is issued if the monitor fast timeout thread has been blocked for an unacceptable amount of time. The event indicates a heavy load on the system and may result in an unexpected (false) takeover.

### Corrective Action

Check CPU load and make sure system is not over subscribed.

### Syslog Message

WARNING failover monitor fast timeout was blocked for %lld secs

## Parameters

**secs** (LONGINT): Number of seconds that the High Availability (HA) node has been blocked

## cf.fm.gbCancelledDuetoDR

### Severity

ERROR

### Description

This event is issued when a giveback has been cancelled due to an ongoing metrocluster disaster recovery operation.

### Corrective Action

Check the status of metrocluster disaster recovery operation by executing command 'metrocluster operation show'. If the command reports metrocluster disaster recovery operation is in progress wait for it to complete and then issue a manual giveback.

### Syslog Message

Failover monitor: giveback cancelled

### Parameters

(None).

## cf.fm.givebackCancelled

### Severity

NOTICE

### Description

This message occurs when a giveback is canceled due to a preexisting state, such as an active CIFS session, a reconstruction, and so on.

### Corrective Action

To override, use the "storage failover giveback -override-vetoes true" command.

### Syslog Message

Failover monitor: giveback canceled.

### Parameters

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.givebackComplete

### Severity

NOTICE

### Description

This message occurs when giveback succeeds.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: giveback completed

**Parameters**

**token** (STRING): Unique token that identifies a failover instance.

**partner\_node\_uuid** (STRING): UUID of the partner node.

**cf.fm.givebackDuration****Severity**

NOTICE

**Description**

This message occurs when a giveback is completed successfully.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: giveback duration time is %llu seconds.

**Parameters**

**giveback\_duration** (LONGINT): Giveback duration time.

**cf.fm.givebackFailed****Severity**

ALERT

**Description**

This message occurs when the failover monitor determines that a giveback has failed. The reason code is a string that describes the reason for the failure.

**Corrective Action**

Resolve the issue based on the reason logged in the message.

**Syslog Message**

Failover monitor: giveback failed '%s'

**Parameters**

**reason** (STRING): Internal reason code for the failure.

**token** (STRING): Unique token that identifies a failover instance.

**partner\_node\_uuid** (STRING): UUID of the partner node.

**cf.fm.givebackForced**

## Severity

ALERT

## Description

This message occurs when the takeover node detects that the takeover process has not been completed within the expected time, and/or normal attempts to give back partner resources also fail. Subsequent to this event, the takeover node will panic and reboot.

## Corrective Action

Attempt to find the panic string in the event logs by using the "event log show" command from the CLI, and then look up the string by using the Panic Message Analyzer tool on the NetApp support site: <http://mysupport.netapp.com/NOW/cgi-bin/pmsg/>. Contact NetApp technical support to confirm the analysis.

## Syslog Message

Failover monitor: forcing reboot to clear state.

## Parameters

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.givebackStarted

### Severity

NOTICE

### Description

This message occurs when the failover monitor initiates a giveback.

### Corrective Action

(None).

### Syslog Message

Failover monitor: giveback started with token %s. "override-vetoes" set to %s, and "require-partner-waiting" set to %s.

### Parameters

**token** (STRING): Unique token that identifies a failover instance.

**override\_vetoes** (STRING): Flag that indicates whether the system overrides veto checks during a giveback operation. This flag corresponds to the "-override-vetoes" parameter of the "storage failover giveback" command. When the parameter is set to true, some veto checks made by subsystems on the source node might be overridden.

**require\_partner\_waiting** (STRING): Flag that indicates whether, during a giveback, the storage is given back regardless of whether the partner node is available to take back the storage. This flag corresponds to the "-require-partner-waiting" parameter of the "storage failover giveback" command. When set to true, the parameter might cause the giveback to proceed, even if the destination node is not ready to receive the aggregate being migrated.

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.givebackUpdateFail

### Severity

ALERT

## Description

This message occurs when GIVEBACK\_DONE is not written to the backup mailbox after all other giveback processing is done. The issuing node is no longer in takeover mode, but the partner node cannot boot (without operator intervention) because the partner mailbox claims it has been taken over.

## Corrective Action

Boot the previously taken over node. During the boot operation, the node requests confirmation to proceed.

## Syslog Message

Failover Monitor: Unexpected error %d while trying to update backup mailbox during giveback

## Parameters

**errcode** (INT): Error code.

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.haltUpdateFail

### Severity

INFORMATIONAL

### Description

This event is issued if we are unable to update the partner state as part of halt processing. This occurrence of this event should not affect the operation of the High Availability (HA) pair.

### Corrective Action

(None).

### Syslog Message

halt: Unable to update failover monitor with NoTakeover state

### Parameters

(None).

## cf.fm.hogger

### Severity

ERROR

### Description

This message occurs when the fast timeout thread is blocked for a very long time and the system can identify threads that might have been responsible for the fast timeout thread not being scheduled.

### Corrective Action

Determine why the process is consuming the CPU, and either correct the problem, or end the offending process.

### Syslog Message

Failover monitor: Process %s ran continuously for %llu ms.

### Parameters

**procName** (STRING): Name of the process that is consuming the CPU.

**schedTime** (LONGINT): Time for which the process ran without releasing the CPU.

## **cf.fm.initError**

### **Severity**

ALERT

### **Description**

This message occurs when failover monitor initialization fails. If this event occurs, the failover monitor cannot be started. The node will reboot after this event.

### **Corrective Action**

Check the logs for other messages from the failing component listed in the message by using the "event log show" command from the CLI. Also check for errors from other components or errors indicating hardware failures. If the problem occurs again after the node reboots, contact NetApp technical support.

### **Syslog Message**

Failover monitor: initialize(%s) fails.

### **Parameters**

**component** (STRING): Software component that has failed to initialize.

## **cf.fm.kernelMismatch**

### **Severity**

ERROR

### **Description**

This event is issued when we detect a possible mismatch of kernel versions in the High Availability (HA) pair. This situation is allowed, although takeover may be disabled if the mismatch imposes version differences in the metadata formats (nvram, filesystem, etc.) of the system.

### **Corrective Action**

Upgrade both nodes to the same release.

### **Syslog Message**

Failover monitor: possible kernel mismatch detected local '%s', partner '%s'

### **Parameters**

**myVersion** (STRING): My version

**partnerVersion** (STRING): The partner's version

## **cf.fm.kernelMismatchOk**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when we detect a possible mismatch of kernel versions in the High Availability (HA) pair has been resolved.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: possible kernel mismatch resolved

**Parameters**

(None).

**cf.fm.launch****Severity**

INFORMATIONAL

**Description**

This event is issued when the failover monitor is launched. It occurs very early in the system startup sequence.

**Corrective Action**

(None).

**Syslog Message**

Launching failover monitor

**Parameters**

(None).

**cf.fm.ImgrVetoOverride****Deprecated**

Deprecated as of version 9.7.

**Severity**

NOTICE

**Description**

This message occurs during an SFO aggregate giveback, when system settings indicate that giveback should be vetoed but the veto was overridden by the automated nondisruptive update procedure. The automated nondisruptive update procedure verifies the expected state of aggregate.

**Corrective Action**

(None).

**Syslog Message**

"%s" subsystem veto was overridden during giveback operation of "%s" aggregate.

**Parameters**

**subsystem** (STRING): Name of the vetoed subsystem.

**aggregate** (STRING): Name of the aggregate.



## **cf.fm.localmbReadStatus**

### **Severity**

INFORMATIONAL

### **Description**

This message reports the status of a local mailbox disk read.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**returncode** (INT): Status returned by the read of the local mailbox disk.

## **cf.fm.lowMemory**

### **Severity**

ALERT

### **Description**

This message occurs when the local node does not have sufficient memory to run failover monitor services.

### **Corrective Action**

Verify that the recommended amount of memory is installed on the system. If there is sufficient memory, the error might be related to hardware issues. In this case, capture the console logs, and then call NetApp technical support.

### **Syslog Message**

Takeover is disabled due to insufficient memory.

### **Parameters**

(None).

## **cf.fm.MBstatusOnBoot**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs on system boot when the failover monitor detects that no takeover is in progress.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

**Parameters**

**status** (INT): Failover monitor status as reported by the mailbox disk.

**cf.fm.mirrorConsistencyOff****Severity**

ERROR

**Description**

This message occurs when the system discovers that the NVRAM mirror consistency option has been disabled. This option should ONLY be disabled under operator control. If mirror consistency is disabled, a takeover can result in a loss of recently logged data.

**Corrective Action**

Run the "cf enable mirrorconsistency" advanced privilege nodeshell command to reenable mirror consistency.

**Syslog Message**

Failover monitor: NVRAM mirror consistency is disabled.

**Parameters**

(None).

**cf.fm.missingAdapter****Severity**

ERROR

**Description**

This message occurs when the HA mode is set to "ha" but no interconnect adapter is found. This is an error indicating a misconfiguration of the system.

**Corrective Action**

Install the high-availability (HA) interconnect adapter or set the HA mode to "non\_ha" by using the "storage failover modify -mode non\_ha" command.

**Syslog Message**

Warning: HA mode is set to "ha" but the interconnect adapter was not found.

**Parameters**

(None).

**cf.fm.monitorBlocked****Severity**

ERROR

**Description**

This event is issued if the failover monitor has been blocked for an unacceptable amount of time. The event indicates a heavy load on the system and may result in an unexpected (false) takeover.

**Corrective Action**

Check CPU load and make sure system is not over subscribed.

**Syslog Message**

WARNING failover monitor was blocked for %lld secs

**Parameters**

**secs** (LONGINT): Number of seconds that the failover monitor has been blocked

**cf.fm.noearlyrelease****Severity**

INFORMATIONAL

**Description**

This message occurs when an early release of reservations is not done.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**state** (INT): Partner firmware state.

**version** (INT): Partner firmware version.

**cf.fm.nofwUpdateinTO****Severity**

INFORMATIONAL

**Description**

This message occurs when there is no progress in the firmware status received from the partner.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**cf.fm.noICbutFoundMb****Severity**

INFORMATIONAL

**Description**

This message occurs when no firmware state is obtained over the High Availability (HA) interconnect but the mailbox disks are found.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**status** (INT): Status of the active/active configuration based on the mailbox disks.

**cf.fm.nombdisks****Severity**

INFORMATIONAL

**Description**

This messages indicates the status of the local mailbox disks.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**returncode** (INT): Return value from the call to read the local mailbox disks.

**mbstatus** (INT): Current status of the active/active configuration.

**cf.fm.noMBdisksOnSFUP****Severity**

ERROR

**Description**

This message occurs when no local mailbox disks are detected, even though the partner performed a giveback.

**Corrective Action**

Check connectivity to all disks by running the "run local storage show" command on each partner, and then comparing the results.

**Syslog Message**

Could not find the local mailbox disks after a giveback. Check connectivity to all disks.

**Parameters**

(None).

## cf.fm.noMBDisksOrlc

### Severity

ERROR

### Description

This message occurs when Data ONTAP® cannot access the local mailbox disks and cannot determine partner status through the high-availability (HA) interconnect.

### Corrective Action

Check connectivity to all disks by running the "run local storage show" command on each partner, and then comparing the results. Verify that the interconnect cables are properly cabled.

### Syslog Message

Could not find the local mailbox disks. Could not determine the firmware state of the partner through the HA interconnect.

### Parameters

(None).

## cf.fm.noPartnerVariable

### Severity

ERROR

### Description

This message occurs when the system cannot identify the serial number of the partner because the firmware variable is not set.

### Corrective Action

1) Use the "storage failover show" command to verify that that high-availability (HA) is enabled. 2) If HA is enabled, there might be too many environment variables defined. Halt the system, and then enter the "printenv" command at the LOADER prompt. Use the "unsetenv" command to remove unneeded environment variables.

### Syslog Message

Unknown partner serial number: firmware %s variable is not set.

### Parameters

**variable** (STRING): Name of the firmware variable.

## cf.fm.noTakeoverNoRc

### Severity

ERROR

### Description

This message indicates that we cannot do takeover during a no-rc boot.

### Corrective Action

Reboot the node normally

**Syslog Message**

Failover monitor: reboot normally to enable takeover

**Parameters**

(None).

**cf.fm.notkoverBadMbox****Severity**

NOTICE

**Description**

This event is issued when we discover that a mailbox is uninitialized.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: uninitialized %s mailbox data detected

**Parameters**

**whose** (STRING): Indicates which mailbox is uninitialized

**cf.fm.notkoverClusterDisable****Severity**

ERROR

**Description**

This event is issued when we discover that failover between the High Availability (HA) pair has been disabled. Failover may be disabled under operator control or when a condition has been discovered (e.g., kernel mismatch) that necessitates disabling of the HA pair.

**Corrective Action**

Resolve the reason provided in the message.

**Syslog Message**

Failover monitor: takeover disabled (%s)

**Parameters**

**reason** (STRING): The reason code for disabling the HA pair

**cf.fm.notkoverOperatorDeny****Severity**

ERROR

**Description**

This event is issued when we discover that the operator has disabled takeover-by-partner.

**Corrective Action**

If takeover by the partner is desired, re-enable takeover.

**Syslog Message**

Failover monitor: takeover by partner disabled

**Parameters**

(None).

**cf.fm.notkoverOperatorDisableNvram****Severity**

ERROR

**Description**

This event is issued when we discover that the operator has disabled the nvram mirror.

**Corrective Action**

Re-enable NVRAM mirroring

**Syslog Message**

Failover monitor: nvram mirror disabled

**Parameters**

(None).

**cf.fm.overwriteState****Severity**

NOTICE

**Description**

This event is issued when the operator has manually intervened and has forced an overwrite of failover monitor state.

**Corrective Action**

(None).

**Syslog Message**

System continuing after overwriting failover monitor state!

**Parameters**

(None).

**cf.fm.panicAfterToDone****Severity**

ALERT

## Description

This message occurs when a node panics too soon after the completion of a takeover. The node reboots in normal mode to avoid recursive panics.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Failover monitor: Panic occurred too soon after takeover was completed (currentTime %llu ms, Takeover completed %llu ms).

## Parameters

**currentTime** (LONGINT): Time when the panic occurred.

**ToDoneTime** (LONGINT): Time when the takeover was completed.

## cf.fm.panicInToMode

### Severity

EMERGENCY

### Description

This message occurs when the node panics after taking over the partner node. When the node comes back up, it will do so in takeover mode.

### Corrective Action

Attempt to find the panic string in the event logs by using the "event log show" command from the CLI, and then look up the string by using the Panic Message Analyzer tool on the NetApp support site: <http://mysupport.netapp.com/NOW/cgi-bin/pmsg/>. Contact NetApp technical support to confirm the analysis.

### Syslog Message

Failover monitor: Panic in takeover mode; takeover will occur on reboot.

### Parameters

(None).

## cf.fm.panicOnGBforced

### Severity

ALERT

### Description

This message occurs when a node panics while a forced giveback is in progress. The node performs giveback and releases partner resources on reboot.

### Corrective Action

Capture the console log and contact NetApp technical support.

### Syslog Message

Failover monitor: Panic during forced giveback; node will release partner resources on reboot.



**Parameters**

(None).

**cf.fm.panicToInProgress****Severity**

ALERT

**Description**

This message occurs when a node panics while the takeover is in progress. The node reboots in normal mode with takeover disabled.

**Corrective Action**

Capture the console log and contact NetApp technical support.

**Syslog Message**

Failover monitor: Panic during takeover; takeover will be disabled on reboot.

**Parameters**

(None).

**cf.fm.partner****Severity**

INFORMATIONAL

**Description**

This event is issued to announce the name of the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner '%s'

**Parameters**

**partner** (STRING): The name of the High Availability (HA) partner

**cf.fm.partnerChange****Severity**

INFORMATIONAL

**Description**

This event is issued to announce a change in the name of the partner.

**Corrective Action**

(None).

### Syslog Message

Failover monitor: partner hostname has changed: '%s'

### Parameters

**partner** (STRING): The name of the High Availability (HA) partner

## cf.fm.partnerFwState

### Severity

INFORMATIONAL

### Description

This message reports the firmware status of the partner.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**state** (INT): Partner firmware status.

## cf.fm.partnerFwTransition

### Severity

INFORMATIONAL

### Description

This message occurs when there is a change in the partner firmware state.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**prevstate** (STRING): Previously reported partner firmware state.

**newstate** (STRING): New firmware state, as reported by the partner.

**progresscounter** (LONGINT): New progress counter, as reported by the partner.

## cf.fm.partnerICFwVersion

### Severity

INFORMATIONAL

### Description

This message occurs when the partner is using a different version of the interconnect firmware.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**version** (INT): Partner firmware version.

**cf.fm.partnerSysid****Severity**

INFORMATIONAL

**Description**

This event is issued to announce the system id of the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner system id: %u

**Parameters**

**sysid** (LONGINT): The sysid of the High Availability (HA) partner

**cf.fm.partnerSysidChange****Severity**

INFORMATIONAL

**Description**

This event is issued to announce a change in the system id of the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner system id has changed: %u

**Parameters**

**sysid** (LONGINT): The sysid of the High Availability (HA) partner

**cf.fm.partnerVolumesOnline****Severity**

NOTICE

**Description**

This event is issued to indicate that the partner's volumes have been brought on-line as part of early takeover processing.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner volumes on-line

**Parameters**

(None).

**cf.fm.replayOnlyTakeover****Severity**

INFORMATIONAL

**Description**

This event is issued when the failover monitor initiates a replay-only takeover, which essentially means performing takeover till the partner logs have been replayed, and then initiating a giveback.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Starting replay-only takeover. A giveback will be initiated after the partner logs have been replayed.

**Parameters**

(None).

**cf.fm.replayOnReboot****Severity**

INFORMATIONAL

**Description**

This message occurs if a node panics in takeover mode and replay of the partner logs will be attempted on reboot.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: replay of partner logs will be attempted on reboot.

**Parameters**

(None).

**cf.fm.reserveDisksOff****Severity**

EMERGENCY

## Description

This event is issued if we discover that disk reservations have been disabled. If this event occurs, contact NetApp technical support.

## Corrective Action

(Call support).

## Syslog Message

Failover monitor: disk reservations disabled

## Parameters

(None).

## cf.fm.reserveMBproblem

### Severity

ERROR

### Description

This message occurs when ONTAP® cannot reserve a high-availability (HA) partner mailbox disk during a takeover.

### Corrective Action

Check connectivity to all disks by using the "storage disk show -fields diskpathnames" command to verify each node in the HA pair has access to all disks. If some disks are not fully accessible, confirm the disks are correctly cabled. To check whether one HA node cannot access disks that are visible to the HA partner node, use the "storage failover show -fields local-missing-disks, partner-missing-disks" command.

### Syslog Message

Takeover has been aborted because the partner mailbox disk: %s could not be reserved. Error: %u.

### Parameters

**diskname** (STRING): Partner mailbox disk that ONTAP could not reserve.

**disk\_error** (INT): Disk reservation error that was encountered.

## cf.fm.slowTimeoutBlocked

### Severity

NOTICE

### Description

This message occurs when the High Availability slow timeout thread has been blocked for an unacceptable amount of time. The event indicates a heavy load on the system and may result in an unexpected takeover.

### Corrective Action

Check CPU load and make sure system is not over subscribed. Contact NetApp technical support for further assistance.

### Syslog Message

High Availability slow timeout was blocked for %lld secs.

**Parameters**

**secs** (LONGINT): Number of seconds that the High Availability (HA) slow timeout thread has been blocked.

**cf.fm.smsVetoOverride****Deprecated**

Deprecated as of version 9.7.

**Severity**

NOTICE

**Description**

This message occurs during an SFO aggregate giveback, when the SnapMirror® subsystem indicates that giveback should be vetoed but the veto was overridden by the automated nondisruptive update procedure. The automated nondisruptive update procedure verifies the expected state of the aggregate.

**Corrective Action**

(None).

**Syslog Message**

"%s" subsystem veto was overridden during giveback operation of "%s" aggregate.

**Parameters**

**subsystem** (STRING): Name of the vetoed subsystem.

**aggregate** (STRING): Name of the aggregate.

**cf.fm.softError****Severity**

ERROR

**Description**

This event is issued when a "soft error" has occurred in the failover monitor.

**Corrective Action**

Resolve the failure listed in the message.

**Syslog Message**

Failover monitor: %s

**Parameters**

**reason** (STRING): Description of the failure.

**cf.fm.takeoverComplete****Severity**

NOTICE

**Description**

This message occurs when a takeover succeeds.

## Corrective Action

(None).

## Syslog Message

Failover monitor: takeover completed

## Parameters

**token** (STRING): Unique token that identifies a failover instance.

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.takeoverDetectionSeconds.Default

### Severity

ERROR

### Description

This message occurs when the takeover detection time is set to a value less than the `DEFAULT_FIRMWARE_TIMEOUTS` setting. This can result in false takeovers and takeovers without diagnostic core dumps.

### Corrective Action

Modify the takeover detection time to the recommended value by using the "storage failover modify -detection-time" command.

### Syslog Message

Takeover detection time is set to %d seconds, which is below the recommended value of %d seconds. False takeovers and takeovers without diagnostic core dumps might occur.

### Parameters

**SECONDS** (INT): Value that the takeover detection time is set to.

**FIRMWARE\_TIMEOUT\_DEF** (INT): Recommended value.

## cf.fm.takeoverDetectionSeconds.Kernel

### Severity

ERROR

### Description

This message occurs when the takeover detection time is set to a value less than the `KERNEL_TIMEOUT` setting (as specified by the "sk.process.timeout.override" option). This can result in takeovers without accompanying diagnostic core dumps of the taken over node.

### Corrective Action

Set the takeover detection time to the recommended value by using the "storage failover modify -detection-time" command.

### Syslog Message

Takeover detection time is set to %d seconds, which is below %d (= sk.process.timeout.override + 5) seconds. Takeovers without diagnostic core dumps might occur.

## Parameters

**SECONDS** (INT): Value that the takeover detection time is set to.

**KERNEL\_TIMEOUT** (INT): Minimum value that should be used.

## cf.fm.takeoverDuration

### Severity

INFORMATIONAL

### Description

This message occurs when a takeover is completed successfully.

### Corrective Action

(None).

### Syslog Message

Failover monitor: takeover duration time is %llu seconds.

## Parameters

**takeover\_duration** (LONGINT): Takeover duration time.

## cf.fm.takeoverFailed

### Severity

ALERT

### Description

This message occurs when the failover monitor determines that a takeover has failed. The reason code is a string that describes the reason for the failure. Any data LIFs that were migrated as part of the takeover operation are not automatically reverted.

### Corrective Action

Resolve the issue based on the reason logged in the message.

### Syslog Message

Failover monitor: takeover failed '%s'

## Parameters

**reason** (STRING): Internal reason code for the failure.

**token** (STRING): Unique token that identifies a failover instance.

**partner\_node\_uuid** (STRING): UUID of the partner node.

## cf.fm.takeoverStarted

### Severity

NOTICE

### Description

This message occurs when the failover monitor initiates a takeover.



**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover started

**Parameters**

**token** (STRING): Unique token that identifies a failover instance.

**partner\_node\_uuid** (STRING): UUID of the partner node.

**cf.fm.timeMasterStatus****Severity**

INFORMATIONAL

**Description**

This event is when we determine our status as time master or slave.

**Corrective Action**

(None).

**Syslog Message**

Acting as time %s

**Parameters**

**masterOrSlave** (STRING): Master or Slave

**cf.fm.TODetectionSecs.reset****Severity**

INFORMATIONAL

**Description**

This message occurs when the current setting of takeover detection time is shorter than the minimum takeover detection time allowed by this version of Data ONTAP®. This can result in false takeovers or takeovers without diagnostic core dumps. Data ONTAP resets the takeover detection time to the new minimum.

**Corrective Action**

(None).

**Syslog Message**

Takeover detection time was set to %d seconds, shorter than the minimum allowed. Reset the detection time to a new minimum of %d seconds.

**Parameters**

**SECONDS** (INT): Current takeover detection seconds.

**FIRMWARE\_TIMEOUT\_DEF** (INT): New default takeover detection seconds.

## cf.fm.transitTimeChange

### Severity

INFORMATIONAL

### Description

This message occurs when you set the takeover or giveback transit timeout to a value other than the default value. During takeover or giveback, if the timeout is exceeded by a subsystem during the takeover/giveback processing, a panic occurs. If the timeout is set too high, longer client outages might occur instead of aborting the takeover/giveback.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**SECONDS** (INT): Transit timeout value (in seconds).

**DEFAULT\_VAL** (INT): Default transit timeout value (in seconds).

## cf.fm.undoFailedTakeover

### Severity

NOTICE

### Description

This event is issued when we initiate an undo of a failed takeover.

### Corrective Action

(None).

### Syslog Message

Failover monitor: initiate giveback due to failed takeover

### Parameters

(None).

## cf.fm.unexpectedPartner

### Severity

ERROR

### Description

This message occurs when the HA mode is set to "non\_ha" but the HA mode was set to "ha" previously. This is not an error, but indicates a possible misconfiguration of the system.

### Corrective Action

Determine whether the HA mode should be set to "ha", and if so, set it.

## Syslog Message

Warning: HA mode is set to "non\_ha" but the node once had a storage failover partner.

## Parameters

(None).

## cf.fm.versionMismatch

### Severity

ALERT

### Description

This event occurs when a version mismatch is detected during internode initialization. Each node transmits its version information to its partner. If a mismatch is detected, the High Availability (HA) takeover capability is disabled.

### Corrective Action

Boot both nodes on the same release.

## Syslog Message

Failover monitor: %s version mismatch detected: %d/%d

## Parameters

**subsystem** (STRING): The name of the versioned subsystem

**myVersion** (INT): My version

**partnerVersion** (INT): The partner's version

## cf.fm.waitBeforeWFG

### Severity

INFORMATIONAL

### Description

This message occurs when a system waits, during boot, for a module to come up before declaring itself ready for giveback. Examples include waiting for the NVRAM battery to be charged.

### Corrective Action

(None).

## Syslog Message

Failover monitor: waited %llu seconds for module %s.

## Parameters

**secs** (LONGINT): Amount of time spent waiting, in seconds.

**module\_name** (STRING): Name of the module the system is waiting for.

## cf.fmmbx events

## cf.fmmbx.mccip.policy.enabled

### Severity

NOTICE

### Description

This message occurs when the system has enabled MCCIP mailbox disk IO policy

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

(None).

## cf.fmns events

### cf.fmns.skipped.disk

### Severity

NOTICE

### Description

This message occurs when a node is in "Waiting For Giveback" state and the Failover Monitor Node State (fmns) module skips a local disk while releasing the reservations in that state.

### Corrective Action

(None).

### Syslog Message

While releasing the reservations in "Waiting For Giveback" state Failover Monitor Node State(fmns) module skipped the disk %s that is owned by %u and reserved by %u.

### Parameters

**uid\_string** (STRING): Name of the disk which was skipped by fmns.

**down\_owner** (INT): The system ID of the node that owns this disk.

**resvn\_owner** (INT): The system ID of the node that has a reservation on this disk.

## cf.fsm events

### cf.fsm.autoGBTkoverOnPanic

### Severity

INFORMATIONAL

### Description

This message occurs when a node initiates an automatic giveback to its partner following a takeover that was due to a panic on the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Automatic giveback was initiated following a takeover that was caused by a panic on the partner.

**Parameters**

(None).

**cf.fsm.autoGBTkoverOnReboot****Severity**

INFORMATIONAL

**Description**

This message occurs when a node initiates an automatic giveback to its partner following a takeover that was due to a reboot of the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Automatic giveback was initiated following a takeover that was caused by the partner reboot.

**Parameters**

(None).

**cf.fsm.autoGivebackAttemptsExceeded****Severity**

EMERGENCY

**Description**

This event is issued when auto giveback is disabled due to a ping-pong situation (autoGB followed by t/o followed by autoGB...)

**Corrective Action**

Examine the logs and/or console output from the partner node. Resolve the issue which prevents the node from staying up.

**Syslog Message**

Failover monitor: Automatic giveback is being disabled due to exceeding %d attempts in %d minutes.

**Parameters**

**attempts** (INT): Number of automatic givebacks attempted

**minutes** (INT): Time period where the automatic givebacks were attempted

## cf.fsm.autoGivebackDelayed

### Severity

INFORMATIONAL

### Description

This message occurs when an automatic giveback is delayed because 'Delay Before Auto Giveback' is set to a non-zero number. If you want to eliminate the delay before automatic giveback occurs, you can use the command "storage failover modify -delay-seconds" to set it to zero.

### Corrective Action

(None).

### Syslog Message

Failover monitor: Automatic giveback was delayed by %d seconds due to a non-zero value of 'Delay Before Auto Giveback'.

### Parameters

**seconds** (INT): Number of seconds by which automatic giveback was delayed.

## cf.fsm.autoGivebackReleasingReservations

### Severity

NOTICE

### Description

This event is generated when we release the disk reservations in preparation for an automatic giveback.

### Corrective Action

(None).

### Syslog Message

Failover monitor: Releasing disk reservations in preparation for an automatic giveback

### Parameters

(None).

## cf.fsm.autoGivebackStarted

### Severity

INFORMATIONAL

### Description

This event is issued when an automatic giveback is initiated.

### Corrective Action

(None).

### Syslog Message

Failover monitor: Automatic giveback started

**Parameters**

(None).

**cf.fsm.autoGivebackVetoed****Severity**

ERROR

**Description**

This event is issued when one or more subsystems have vetoed the automatic giveback.

**Corrective Action**

Terminate the long-running jobs and auto giveback will be successful next time it is attempted.

**Syslog Message**

Failover monitor: Automatic giveback has been deferred due to long running operations

**Parameters**

(None).

**cf.fsm.autoTakeoverCancelled****Severity**

NOTICE

**Description**

This event is issued when the failover monitor cancels a pending takeover.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: pending takeover cancelled

**Parameters**

(None).

**cf.fsm.backupMailboxError****Severity**

ERROR

**Description**

This message occurs when the failover monitor determines that an error was observed in the partner's mailbox.

**Corrective Action**

Correct the issues preventing the node from accessing the partner's mailbox disks. Check for cabling, host bus adapter (HBA), storage controller or drive/LUN issues. You can also use Multipathing to provide a redundant connection to the mailbox disk.

**Syslog Message**

Failover monitor: partner mailbox error detected.

**Parameters**

(None).

**cf.fsm.backupMailboxOk****Severity**

NOTICE

**Description**

This event is issued when the failover monitor has determined that an error in the backup's mailbox has been fixed.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: backup mailbox OK

**Parameters**

(None).

**cf.fsm.checkingForBigJobs****Severity**

INFORMATIONAL

**Description**

This event is generated when we want to initiate an automatic giveback and we're checking for long running operations which might veto our plans.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Checking for long running operations in preparation for an automatic giveback.

**Parameters**

(None).

**cf.fsm.cliTakeoverCancelled****Severity**

NOTICE

**Description**

This event is issued when the failover monitor cancels a pending takeover issued through a CLI.



**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover cannot be performed because of reason (%s)

**Parameters**

**reason** (STRING): Reason why takeover cannot occur

**cf.fsm.firmwareStatus****Severity**

INFORMATIONAL

**Description**

This event is issued when we detect an altered firmware status update from the partner.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner %s

**Parameters**

**reason** (STRING): Partner status

**cf.fsm.givebackHung****Severity**

ALERT

**Description**

This event occurs when the failover monitor detects that the giveback process is hung.

**Corrective Action**

Collect the resulting core file and provide it to Customer Support.

**Syslog Message**

Failover monitor: giveback process is hung ('%s')

**Parameters**

**moduleName** (STRING): The name of the module that the hang occurred in.

**cf.fsm.givebackMaxRetries****Severity**

ALERT

**Description**

This event is called when the giveback retry count has been exceeded. This situation exists when the system is unable either to takeover or giveback. It may be due to either a hardware bug (e.g., the disk

subsystem is hung) or a software bug.

#### **Corrective Action**

Examine the logs and determine why the giveback is failing. Correct that problem and retry the giveback.

#### **Syslog Message**

Failover monitor: giveback has exceeded max retry count

#### **Parameters**

**retries** (INT): Number of retries attempted.

### **cf.fsm.nfo.acceptTakeoverReq**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the system starts a negotiated takeover of its partner, and requests a graceful shutdown of the partner.

#### **Corrective Action**

(None).

#### **Syslog Message**

Negotiated failover: starting takeover and shutdown of partner (%s), will take over in at most %d secs.  
Reason: %s.

#### **Parameters**

**partnerName** (STRING): Name of partner node.

**maxTakeoverTime** (INT): Maximum amount of time to wait for the partner to shut down before starting takeover, in seconds.

**partnerReason** (STRING): Reason for the initiation of the takeover.

### **cf.fsm.nfo.clearReqTakeover**

#### **Severity**

NOTICE

#### **Description**

This event is called when the system clears a request for takeover by its partner.

#### **Corrective Action**

(None).

#### **Syslog Message**

Negotiated failover: clearing partner takeover request

#### **Parameters**

(None).

## **cf.fsm.nfo.delayingGracefulShutdown**

### **Severity**

ERROR

### **Description**

This event is called when the system has been asked to shutdown by its partner as the result of the negotiated failover mechanism, but the system can not shut down due to a specific reason.

### **Corrective Action**

Using the information provided in the message, determine why shutdown cannot be invoked. Resolve that problem and retry the takeover request.

### **Syslog Message**

Negotiated failover: delaying shutdown due to %s

### **Parameters**

**why** (STRING): Indicates the cause of the delay.

## **cf.fsm.nfo.disable**

### **Severity**

ERROR

### **Description**

This event is called when negotiated failover is disabled for a particular module.

### **Corrective Action**

Examine previous messages for failures related to the type of NFO.

### **Syslog Message**

Negotiated failover: disabling negotiated failover for module %s

### **Parameters**

**mod** (STRING): Negotiated failover module or type.

## **cf.fsm.nfo.disable.shelfversion**

### **Severity**

ERROR

### **Description**

This event is called when negotiated failover is disabled due to Shelf Count message version mismatch.

### **Corrective Action**

Upgrade both nodes to the same release.

### **Syslog Message**

Negotiated failover: disabling negotiated failover due to version mis-match.

**Parameters**

(None).

**cf.fsm.nfo.enable****Severity**

NOTICE

**Description**

This event is called when negotiated failover is enabled for a particular module.

**Corrective Action**

(None).

**Syslog Message**

Negotiated failover: enabling negotiated failover for module %s

**Parameters**

**mod** (STRING): NFO module (or type)

**cf.fsm.nfo.gracefulShutdownTimerExpired****Severity**

NOTICE

**Description**

This event is called when the maximum time the system will wait for the partner to shutdown gracefully has passed. At this point the system takes over by force.

**Corrective Action**

(None).

**Syslog Message**

Negotiated failover: partner graceful shutdown appears hung, taking over

**Parameters**

(None).

**cf.fsm.nfo.moduleImpairmentChange****Severity**

ERROR

**Description**

This event is called when a module which is participating in negotiated failover changes from "unimpaired" to "impaired" or vice versa.

**Corrective Action**

Check the state of the module listed in the message.

## Syslog Message

Negotiated failover: module %s is now %s

## Parameters

**mod** (STRING): Type of negotiated failover

**impairment** (STRING): Either unimpaired or impaired

## cf.fsm.nfo.partnerShutdown

### Severity

NOTICE

### Description

This event is called when the system sees that the partner has finished shutting down gracefully during negotiated failover.

### Corrective Action

(None).

## Syslog Message

Negotiated failover: partner has shutdown

## Parameters

(None).

## cf.fsm.nfo.rejectTakeoverReqDueToImpairment

### Severity

ERROR

### Description

This message occurs when the system rejects a request by its partner to take it over because the system is itself impaired.

### Corrective Action

Use the "storage failover show" command to determine why takeover is not possible. Resolve that problem and retry the takeover request.

## Syslog Message

Negotiated failover: rejecting takeover request by partner due to own impairment.

## Parameters

(None).

## cf.fsm.nfo.rejectTakeoverReqDueToOwnReq

### Severity

ERROR

## Description

This event is called when the system rejects a request by its partner to take it over because the system has itself recently requested takeover by its partner. Rejecting this request prevents each system trying to takeover its partner simultaneously. If the partner persists and the system doesn't become impaired the request will soon be granted.

## Corrective Action

Resolve any impairment issues reported in previous messages that would cause takeover. If takeover is requested by the operator, it should only be requested on one node.

## Syslog Message

Negotiated failover: rejecting takeover request by partner due to own recent takeover request.

## Parameters

(None).

## cf.fsm.nfo.startingGracefulShutdown

### Severity

NOTICE

### Description

This event is called when the system has been asked to shutdown by its partner as the result of the negotiated failover mechanism. The system responds by shutting down gracefully, shutting down services in an orderly manner.

### Corrective Action

(None).

### Syslog Message

Negotiated failover: starting graceful shutdown.

### Parameters

(None).

## cf.fsm.nfo.takeoverPendingCancel

### Severity

ERROR

### Description

This message occurs when the system is waiting for the partner to shutdown gracefully and failover is disabled, canceling the pending takeover.

### Corrective Action

Use the "storage failover modify -enabled true" command to reenable failover.

### Syslog Message

Negotiated failover: pending takeover canceled.

**Parameters**

(None).

**cf.fsm.partnerNotResponding****Severity**

NOTICE

**Description**

This event is issued when we detect that the partner node is not responsive.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner not responding

**Parameters**

(None).

**cf.fsm.partnerOk****Severity**

NOTICE

**Description**

This event is issued when we detect that the partner node, which was previously not responsive, is now OK.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: partner ok

**Parameters**

(None).

**cf.fsm.releasingReservations****Severity**

INFORMATIONAL

**Description**

This event is generated when we release the disk reservations in preparation for a manual giveback.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Releasing disk reservations in preparation for giveback

## Parameters

(None).

## cf.fsm.shelfCount.fewerShelves

### Severity

ALERT

### Description

This event is emitted when we detect that the partner sees more of our disk shelves than we do. In other words, it sees more shelves on its FCAL B loop than we see on our A loop. This is probably due to a cabling problem or a broken FCAL host adaptor. If "disk\_shelf" negotiated failover is enabled, this condition should lead to a takeover by the partner if the partner is otherwise able to take us over.

### Corrective Action

Resolve cabling issues which are preventing both nodes from seeing the same disks.

### Syslog Message

Disk shelf count mismatch: partner sees more of our A shelves on its B loop (%d) than we do (%d).

### Parameters

**bShelves** (INT): Number of our shelves which the partner can see

**aShelves** (INT): Number of shelves which we can see

## cf.fsm.stateTransit

### Severity

INFORMATIONAL

### Description

This event is issued when a state transition is detected. Typically, this indication means that the failover monitor is about to either takeover its partner or giveback to its partner. This can happen as the result of timers going off, operator command, or an indication from the partner that a fault has been detected.

### Corrective Action

(None).

### Syslog Message

Failover monitor: %s → %s

### Parameters

**oldState** (STRING): The old failover monitor state.

**newState** (STRING): The new failover monitor state.

**elem** (STRING): The name of the FSM element that has caused the state transition to occur. This value is dependent upon the FSM implementation.

## cf.fsm.takeover.bootingExp

### Severity

ALERT



**Description**

This event is issued when an automatic takeover is initiated after detecting that the partner boot process is hung trying to load the kernel.

**Corrective Action**

Please capture console log of partner filer and contact Customer Support

**Syslog Message**

Failover monitor: automatic takeover attempted after detecting that partner is hung loading kernel while booting

**Parameters**

(None).

**cf.fsm.takeover.disaster****Severity**

INFORMATIONAL

**Description**

This message occurs when an operator-requested disaster recovery (DR) takeover is initiated.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover attempted after "cf forcetakeover -d" command.

**Parameters**

(None).

**cf.fsm.takeover.early****Severity**

INFORMATIONAL

**Description**

This event is issued when a filer takes over its partner while booting up in takeover mode.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover resumption attempted after reboot

**Parameters**

(None).

**cf.fsm.takeover.forced**

**Severity**

INFORMATIONAL

**Description**

This message occurs when an operator-requested forced takeover is initiated.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover attempted after "cf forcetakeover" or "storage failover takeover -option force" in advanced privilege.

**Parameters**

(None).

**cf.fsm.takeover.mdp****Severity**

ALERT

**Description**

This event is issued when an automatic takeover is initiated after detecting that the partner has panicked due to a multi-disk failure

**Corrective Action**

Please check the connectivity of the partner to it's disks and shelves and contact customer support.

**Syslog Message**

Failover monitor: takeover attempted after multi-disk failure on partner

**Parameters**

(None).

**cf.fsm.takeover.ndu****Severity**

INFORMATIONAL

**Description**

This message occurs when an operator-requested takeover is initiated with the "cf takeover -n" or "storage failover takeover -option allow-version-mismatch" command.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover attempted after "cf takeover -n" or "storage failover takeover -option allow-version-mismatch" command.

**Parameters**

(None).

**cf.fsm.takeover.nfo****Severity**

INFORMATIONAL

**Description**

This message occurs when an operator-requested takeover is initiated with the "storage failover takeover" command.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover attempted after "storage failover takeover" command.

**Parameters**

(None).

**cf.fsm.takeover.noHeartbeat****Severity**

ALERT

**Description**

This message occurs when a node detects no heartbeat from the partner, indicating that the partner is not functioning. The node will attempt an automatic takeover.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Failover monitor: Takeover initiated after no heartbeat was detected from the partner node.

**Parameters**

(None).

**cf.fsm.takeover.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when an operator-requested takeover is initiated with the "storage failover takeover -option immediate" command, or when Kernel Cluster Services Connectivity, Liveness and Availability Monitor (CLAM) triggers a takeover after determining that the partner node is out of "CLAM quorum".

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: %s attempted

**Parameters**

**reason** (STRING): Reason a takeover of the partner was triggered. Possible values are "Operator initiated immediate takeover" and "CLAM initiated takeover".

**cf.fsm.takeover.on.halt****Severity**

INFORMATIONAL

**Description**

This message occurs when one node in a high-availability (HA) pair initiates an automatic takeover after detecting that its partner node has halted.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Node initiated automatic takeover after detecting that its partner node has halted.

**Parameters**

(None).

**cf.fsm.takeover.on.reboot****Severity**

INFORMATIONAL

**Description**

This message occurs when one node in a High Availability (HA) pair initiates an automatic takeover after detecting that its partner node is rebooting.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: One node initiated automatic takeover after detecting that its partner node is rebooting.

**Parameters**

(None).

**cf.fsm.takeover.operatorExp****Severity**

NOTICE

**Description**

This event is issued when an automatic takeover is initiated after detecting that the operator timer has expired. This may happen when the operator has failed to respond to a question or not entered a required command during boot.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover attempted after operator timeout expired on partner

**Parameters**

(None).

**cf.fsm.takeover.panic****Severity**

ALERT

**Description**

This message occurs when an automatic takeover is initiated after detecting that the partner has panicked.

**Corrective Action**

Attempt to find the panic string in the event logs by using the "event log show" command from the CLI, and then look up the string by using the Panic Message Analyzer tool on the NetApp support site: <http://mysupport.netapp.com/NOW/cgi-bin/pmsg/>. Contact NetApp technical support to confirm the analysis.

**Syslog Message**

Failover monitor: takeover attempted after partner panic.

**Parameters**

(None).

**cf.fsm.takeover.postFailed****Severity**

ALERT

**Description**

This event is issued when an automatic takeover is initiated after detecting that the partner's power-on self-test has failed.

**Corrective Action**

Please run hardware diagnostics on partner filer and contact Customer Support

**Syslog Message**

Failover monitor: takeover attempted after partner POST failed

**Parameters**

(None).

## **cf.fsm.takeover.rebootExp**

### **Severity**

ALERT

### **Description**

This message occurs when an automatic takeover is initiated after detecting that the partner boot process is hung.

### **Corrective Action**

Capture console log of partner node and contact NetApp technical support.

### **Syslog Message**

Failover monitor: automatic takeover attempted after detecting that partner is hung in boot.

### **Parameters**

(None).

## **cf.fsm.takeover.shortUptime**

### **Severity**

ALERT

### **Description**

This message occurs when an automatic takeover is initiated after detecting that the partner died very shortly after booting up.

### **Corrective Action**

Contact NetApp technical support.

### **Syslog Message**

Failover monitor: takeover attempted after partner went down shortly after booting up

### **Parameters**

(None).

## **cf.fsm.takeover.sparecoreExp**

### **Severity**

ALERT

### **Description**

This message occurs when an automatic takeover is initiated by the local node after it detects that the partner node has panicked, and that the partner has not initiated a main memory core dump to disk in a reasonable amount of time.

### **Corrective Action**

Capture the console log of the partner node, and then contact NetApp technical support.

### **Syslog Message**

A takeover was attempted by the local node after sparecore timeout expired on the partner node.

## Parameters

(None).

## cf.fsm.takeoverByPartnerDisabled

### Severity

ERROR

### Description

This message occurs when the failover monitor determines that takeover by the partner is disabled.

### Corrective Action

Find the reason for the error message (it is surrounded by parentheses). Based on that reason, the corrective action is one of the following:

- allowed: Takeover is allowed; you do not need to take any action.
- Controller failover (CFO) is not initialized: Make sure that the high-availability (HA) pair is set up correctly. Contact NetApp technical support if you need assistance.
- Controller is in non-HA mode: Set the HA mode to "ha" by using the "storage failover modify -mode ha" command to activate HA functionality.
- Takeover disabled: Use the "storage failover modify -enabled true" command to reenable HA functionality.
- partner mailbox disks not accessible or invalid: Check connectivity to all disks by running the "run local storage show" command on each node, and then comparing the results. Resolve the differences in disks visible to both systems. Verify that the interconnect cables are properly cabled. Failover monitor version mismatch: Make sure that both the local and partner node are running the same version of Data ONTAP®.
- Takeover disabled by partner: Use the "storage failover modify -enabled true" command to reenable HA functionality.
- Takeover disabled by operator: Use the "storage failover modify -enabled true" command to reenable HA functionality.
- NVRAM size mismatch: Make sure that the local NVRAM (nonvolatile random-access memory) size matches the partner node.
- version mismatch: Make sure that both the local and partner nodes are running the same version of Data ONTAP.
- interconnect error: Make sure that the interconnect link is connected and functioning.
- partner booting: Wait for the partner node to complete its booting process, and then try takeover.
- shelf too hot: Make sure that the disk shelf temperature is properly regulated.
- partner is performing revert: Wait for the partner node to complete the revert process, and then try takeover.
- revert is in progress: Wait for the local node to complete the revert process, and then try takeover.
- partner is attempting takeover: Cannot perform a takeover operation while the partner node is attempting a takeover.
- takeover is in progress: The local node is already taken over or is trying to take over the partner node.
- partner halted in notakeover mode: The partner node was most likely halted using the "halt -f" command; reboot the partner node, and then try again.
- unsynchronized log: Make sure that the interconnect link is connected and functioning.
- unknown notakeover reasons: Contact NetApp technical support.
- waiting for partner to recover: The partner has not booted completely after giveback; wait for the partner to come back up completely.
- low memory: Contact NetApp technical support to upgrade.
- local halt in progress: The local node is about to halt; try again after the reboot.
- status of backup mailbox is uncertain: Check connectivity to all disks by running the "run local storage show" command on each node, and then comparing the results. Resolve the differences in disks visible to both systems. Verify that the interconnect cables are properly cabled.
- automatic takeover disabled: Use the "storage failover takeover" command manually.
- metrocluster disaster recovery operation is in progress: The local node is performing a MetroCluster(tm) disaster recovery operation; wait for it to finish, and then try again.
- This node or partner node is in switchover state and the MetroCluster configuration option "node-object-limit" is off in the disaster recovery(DR) group of this node: Retry takeover after doing a switchover.

### Syslog Message

Failover monitor: takeover of %s by %s disabled (%s).

### Parameters

**local** (STRING): Name of the local node.

**partner** (STRING): Name of the partner node.

**reason** (STRING): Reason takeover is disabled.

## **cf.fsm.takeoverByPartnerEnabled**

### **Severity**

NOTICE

### **Description**

This event is issued when the failover monitor determines that takeover by the partner has been enabled.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: takeover of %s by %s enabled

### **Parameters**

**local** (STRING): Name of local node

**partner** (STRING): Name of partner node

## **cf.fsm.takeoverCountdown**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued as part of the takeover countdown processing in the FSM.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: takeover scheduled in %d seconds

### **Parameters**

**secsTillTakeover** (INT): Number of seconds until takeover occurs

## **cf.fsm.takeoverDelayed**

### **Severity**

NOTICE

### **Description**

This event is issued when we are delaying takeover due to status indications received from the partner.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: takeover delayed, partner %s



## Parameters

**reason** (STRING): Description of why takeover is being delayed  
**secsTillTakeover** (INT): Number of seconds until takeover will be started

## cf.fsm.takeoverHung

### Severity

EMERGENCY

### Description

This message occurs when the failover monitor detects that the takeover process is hung. Subsequent to this event, the takeover node will panic.

### Corrective Action

Attempt to find the panic string in the event logs by using the "event log show" command from the CLI, and then look up the string by using the Panic Message Analyzer tool on the NetApp support site: <http://mysupport.netapp.com/NOW/cgi-bin/pmsg/>. Contact NetApp technical support to confirm the analysis.

### Syslog Message

Failover monitor: takeover process is hung ('%s').

## Parameters

**moduleName** (STRING): Name of the module that the hang occurred in.

## cf.fsm.takeoverOfPartnerDisabled

### Severity

ERROR

### Description

This message occurs when the failover monitor determines that takeover of the partner is disabled.

### Corrective Action

Find the reason for the error message (it is surrounded by parentheses). Based on that reason, the corrective action is one of the following: -allowed: Takeover is allowed; you do not need to take any action. -Controller failover (CFO) is not initialized: Make sure that the high-availability (HA) pair is set up correctly. Contact NetApp technical support for assistance. -Controller is in non-HA mode: Set HA mode to "ha" by using the "storage failover modify -mode ha" command to activate HA functionality. -HA takeover disabled: Use the "storage failover modify -enabled true" command to reenab HA functionality. -partner mailbox disks not accessible or invalid: Check connectivity to all disks by running the "storage show" nodeshell command on each node and comparing the results. Resolve the differences in disks visible to both systems. Verify that the interconnect cables are properly cabled. -failover monitor version mismatch: Make sure that both the local and partner node are running the same version of Data ONTAP®. -Takeover disabled by partner: Use the "storage failover modify -enabled true" command to reenab HA functionality. -Takeover disabled by operator: Use the "storage failover modify -enabled true" command to reenab HA functionality. -NVRAM size mismatch: Make sure that the local NVRAM (nonvolatile random-access memory) size matches the partner node. -version mismatch: Make sure that both the local and partner nodes are running the same version of Data ONTAP. -interconnect error: Make sure that the interconnect link is connected and functioning. -partner booting: Wait for the partner node to complete its booting process, and then try takeover. -shelf too hot: Make sure that the disk shelf temperature is properly regulated. -partner is performing revert: Wait for the partner node to complete the revert process, and then try takeover. -revert is in progress: Wait for the local node to complete the revert process, and then try

takeover. -partner is attempting takeover: Cannot perform a takeover operation while the partner node is attempting a takeover. -takeover is in progress: The local node is already taken over or is trying to take over the partner node. -partner halted in notakeover mode: The partner node was most likely halted using the "halt -f" or "system node halt -inhibit-takeover" command; reboot the partner node, and then try again. -unsynchronized log: Make sure that the interconnect link is connected and functioning. -unknown notakeover reasons: Contact NetApp technical support. -waiting for partner to recover: The partner has not booted completely after giveback; wait for the partner to come back up completely. -low memory: Contact NetApp technical support to upgrade. -local halt in progress: The local node is about to halt; try again after the reboot. -status of backup mailbox is uncertain: Check connectivity to all disks by running the "run local storage show" command on each node, and then comparing the results. Resolve the differences in disks visible to both systems. Verify that the interconnect cables are properly cabled. -automatic takeover disabled: Use the "storage failover takeover" command manually. -metrocluster disaster recovery operation is in progress: The local node is performing a MetroCluster(tm) disaster recovery operation; wait for it to finish, and then try again. -This node or the partner node is in switchover state and the MetroCluster configuration option "node-object-limit" is off in the disaster recovery(DR) group of this node: Retry takeover after doing a switchback.

### Syslog Message

Failover monitor: takeover of %s disabled (%s).

### Parameters

**partner** (STRING): Name of the partner node.  
**reason** (STRING): Description of why takeover cannot occur.

## cf.fsm.takeoverOfPartnerEnabled

### Severity

NOTICE

### Description

This event is issued when the failover monitor determines that takeover of the partner has been enabled.

### Corrective Action

(None).

### Syslog Message

Failover monitor: takeover of %s enabled

### Parameters

**partner** (STRING): Name of partner node

## cf.giveback events

### cf.giveback.delayed.cleanup

### Severity

INFORMATIONAL

### Description

This message occurs when giveback of a partner node is delayed due to the cleanup of a previous giveback event.

## Corrective Action

(None).

## Syslog Message

Failover monitor: giveback process was delayed because of a previous giveback event.

## Parameters

(None).

## cf.giveback.disk.check.fail

### Severity

ALERT

### Description

This message occurs when the giveback failed due to a disk inventory check failure.

### Corrective Action

If the giveback failed or is disabled because the node did not receive the disk inventory information from the partner, then retry giveback after a few minutes. If the giveback failed or is disabled because the partner is missing its disks, make sure that the partner has connectivity to all of its disks. To check the list of disks that the partner can see, boot the partner in Maintenance mode and use the "disk show" command or "storage show disk" command to list its disks. To check for missing disks, use the "storage failover show -instance" command. If you want to override the disk inventory check, use the "cf disable diskInventory" command from the advanced privilege level of the nodeshell on the takeover node. However, disabling the disk inventory check can lead to client outage if the partner cannot see all of its disks.

### Syslog Message

cf giveback failed: %s.

### Parameters

**reason** (STRING): Reason for the disk inventory failure.

## cf.giveback.keepingPartnerAggr

### Severity

ERROR

### Description

This message occurs during the giveback process following a failed takeover attempt, when an aggregate belonging to the partner node is not given back to the partner node. The reason is that the aggregate was already under the control of the node performing the takeover before the takeover attempt. Only aggregates that are acquired during the failed takeover attempt are given back.

### Corrective Action

Give back the aggregate to the partner node by using the "storage failover giveback" command.

### Syslog Message

Not giving back aggregate '%s' because it was under the control of this node prior to the failed takeover attempt.

## Parameters

**aggr** (STRING): Name of the aggregate that was not given back.

# cf.headswap events

## cf.headswap.abort

### Severity

INFORMATIONAL

### Description

This message occurs when the headswap on the partner node is aborted.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**partner\_sysid** (LONGINT): Value of the partner system ID after aborting the headswap.

## cf.headswap.clear.partner

### Severity

INFORMATIONAL

### Description

This message occurs when the disks are manually reassigned to a new partner. The head swap information has been cleared.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

(None).

## cf.headswap.complete

### Severity

INFORMATIONAL

### Description

This message occurs when the headswap of a node is completed.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**old\_sysid** (LONGINT): Value of the partner system ID before completion of the headswap.

**new\_sysid** (LONGINT): Value of the partner system ID after completion of the headswap.

## cf.headswap.detected.local

### Severity

INFORMATIONAL

### Description

This message occurs when a headswap on the local node is detected.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**old\_local\_sysid** (LONGINT): System ID of the node that is being replaced.

**new\_local\_sysid** (LONGINT): System ID of the new node that is replacing the old node.

## cf.headswap.detected.partner

### Severity

INFORMATIONAL

### Description

This message occurs when a headswap on the partner node is detected.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**old\_partner\_sysid** (LONGINT): System ID of the partner that is being replaced.

**new\_partner\_sysid** (LONGINT): System ID of the new partner that is replacing the old partner.

## cf.headswap.state.change

### Severity

INFORMATIONAL

### Description

This message occurs when the headswap state changes. The headswap state is written to the primary mailbox. The possible values are HEADSWAP\_NONE, HEADSWAP\_START, HEADSWAP\_CFO\_START,

HEADSWAP\_CFO\_END, and HEADSWAP\_SFO\_START.

#### **Corrective Action**

(None).

#### **Syslog Message**

(None).

#### **Parameters**

**old\_state\_string** (STRING): Old value of the headswap state.

**new\_state\_string** (STRING): New value of the headswap state.

## **cf.hwassist events**

### **cf.hwassist.Default.NodemgmtAddr**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the hardware-assist module automatically picked a local IP address for the hardware-assisted takeover feature.

#### **Corrective Action**

(None).

#### **Syslog Message**

The system automatically chose %s as the local hardware-assist IP address.

#### **Parameters**

**defaultip** (STRING): Default hardware-wassist IP address chosen to receive hardware assist alerts.

### **cf.hwassist.DefaultEnabled**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the system automatically enables the hardware-assist takeover feature. With hardware-assisted takeover, certain failures, such as watchdog reset, POST error, and so on, trigger hardware-assisted takeover through hardware components such as Remote LAN Module (RLM) or Service Processor (SP).

#### **Corrective Action**

(None).

#### **Syslog Message**

Hardware-assisted takeover feature is enabled on this system. The partner can take over this node faster during certain failures.

## Parameters

(None).

## cf.hwassist.DefaultPrtnrAddr

### Severity

NOTICE

### Description

This message occurs when the hardware-assist partner address is not configured and the system automatically picked a partner IP address for the hardware-assisted takeover feature. If a partner node has multiple node management IP addresses, a different partner node management IP address can optionally be set by using the "storage failover modify -hwassist-partner-ip" command. Note: the automatically configured IP address cannot be modified on some platforms.

### Corrective Action

(None).

### Syslog Message

The system automatically chose %s as the hardware-assisted takeover partner address.

## Parameters

**defaultip** (STRING): Default hardware-assist partner IP address chosen to send hardware-assisted takeover alerts.

## cf.hwassist.DefaultPrtnrPort

### Severity

NOTICE

### Description

This message occurs when the partner's port for the hardware-assisted takeover feature is not configured and the system sets the hardware-assist partner port to a default port number. A different partner port number can optionally be set by using the "storage failover modify -hwassist-partner-port" command. Note: the default port number cannot be modified on some platforms.

### Corrective Action

(None).

### Syslog Message

The system set the default hardware-assist partner port number to %d.

## Parameters

**defaultport** (INT): Default hardware-assist partner port number.

## cf.hwassist.empty.MgmtAddr

### Severity

ERROR

**Description**

This message occurs when the local hardware-assist module cannot find at least one node management IP address that has the same capabilities as the partner's hardware components, such as Remote LAN Module (RLM) or Service Processor (SP).

**Corrective Action**

Complete or correct the setup of the Remote LAN Module (RLM) or Service Process (SP) for each node in your cluster.

**Syslog Message**

Cannot find hardware-assist node management IP address for partner node.

**Parameters**

(None).

**cf.hwassist.empty.NodemgmtAddr****Deprecated**

Deprecated as of version 9.1.

**Severity**

ERROR

**Description**

This message occurs when the local hardware-assist module cannot find at least one node management IP address that has the same capabilities as the partner's hardware components, such as Remote LAN Module (RLM) or Service Processor (SP).

**Corrective Action**

Complete or correct the setup of the Remote LAN Module (RLM) or Service Process (SP) for each node in your cluster.

**Syslog Message**

Cannot find hardware-assist node management IP address for partner node.

**Parameters**

(None).

**cf.hwassist.FwUpgrade****Severity**

ERROR

**Description**

This message occurs when hwassist thread detects that partner's hw\_assist hardware firmware does not support hw\_assist feature.

**Corrective Action**

Upgrade the partner hw\_assist hardware firmware to the latest version.



## Syslog Message

HA hw\_assist: Partner(%s) %s firmware does not support the hw\_assist feature.

## Parameters

**partnerName** (STRING): Partner's name.

**hwassist\_hrdwr** (STRING): Type of hw\_assist hardware, for example, Remote LAN module(RLM) or Service Processor(SP).

## cf.hwassist.hwasstActive

### Severity

INFORMATIONAL

### Description

This message occurs when hw\_assist successfully binds to a port to listen for alerts.

### Corrective Action

None.

## Syslog Message

hw\_assist: hw\_assist functionality is active on IP address: %s port: %d

## Parameters

**ipaddress** (STRING): IP address to which bind was attempted.

**port** (INT): Port on which bind was attempted.

## cf.hwassist.invalidConfig

### Severity

NOTICE

### Description

This message occurs when the system encounters invalid hardware-assist configurations.

### Corrective Action

(None).

## Syslog Message

Invalid hardware-assist configurations. IPv4: %s IPv6: %s port: %d keep\_alive\_interval: %d

## Parameters

**ipv4** (STRING): Partner node IPv4 address.

**ipv6** (STRING): Partner node IPv6 address.

**port** (INT): Partner node hardware-assist port number.

**keep\_alive\_interval** (INT): Hardware-assist keep-alive interval.

## cf.hwassist.IpNotConfigured

### Severity

ERROR

## Description

This message occurs when the hwassist thread detects that the partner's hw\_assist hardware is not configured with an IPv4 or IPv6 address.

## Corrective Action

Configure a valid IPv4 or IPv6 address for the partner's hw\_assist hardware by using the "storage failover modify -hwassist-partner-ip" command.

## Syslog Message

cf hw\_assist: Partner(%s)'s %s is not configured with the right IP address family. Use the "storage failover hwassist show" command for details.

## Parameters

**partnerName** (STRING): Partner's name.

**hwassist\_hrdwr** (STRING): Type of hw\_assist hardware, for example, RLM.

## cf.hwassist.keyMismatch

### Severity

NOTICE

### Description

This message occurs when hwassist receives a alert with wrong key. EMS displays both the expected and the received keys.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**expectedKey** (STRING): Expected key, generated by Failover Monitor(FM).

**receivedKey** (STRING): Received key from hwassist hardware.

**systemid** (STRING): Received systemid.

## cf.hwassist.localMonitor

### Severity

ERROR

### Description

This message occurs when hwassist thread is inactive because of problems on local node.

### Corrective Action

check console log for cf.hwassist\_socBindFailed message. If message exist, take corrective action for that message.

### Syslog Message

hw\_assist: hw\_assist functionality is inactive.

## Parameters

(None).

## cf.hwassist.missedKeepAlive

### Severity

ERROR

### Description

This message occurs when the HW-assisted takeover process on this node has not received keep-alive messages from its HA partner.

### Corrective Action

Use the "storage failover hwassist show" command to verify the status of the HW-assisted takeover feature. If the 'Keep Alive Status' field reports "healthy", it is functional. If the "storage failover hwassist show" command "Keep Alive Status" field reports "did not receive hwassist keep alive alerts from partner" and the "storage failover hwassist stats show" command shows the value of the "ID\_mismatch" is incrementing, disable the service-processor (SP or BMC) IPv4 or IPv6 network interfaces and then reenabling them, as stated below. Use the "system service-processor network show" command to display the current network configuration. Note whether the IPv4 or IPv6 address-family is enabled. You need this information in the next steps. Perform the following on each node that has an issue: Disable the interface for an address-family by using the "sp network modify -node {name} -address-family {IPv4|IPv6} -enable false" command. Then, reenabling the interface for an address-family by setting "-enable true" via the above command.

### Syslog Message

HW-assisted takeover missing keep-alive messages from HA partner (%s).

### Parameters

**partnerName** (STRING): HA partner name.

## cf.hwassist.noConfigRecv

### Severity

ERROR

### Description

This message occurs when hwassist thread has not received any config info from partner.

### Corrective Action

None.

### Syslog Message

hw\_assist: hw\_assist functionality inactive. No config received from partner(%s).

### Parameters

**partnerName** (STRING): Partner's name.

## cf.hwassist.notCapable

### Severity

ERROR

## Description

This message occurs when the currently installed firmware version of the hardware module doesn't support the hardware-assisted takeover feature.

## Corrective Action

Upgrade the hardware module firmware to the latest version.

## Syslog Message

Currently installed version of %s's firmware doesn't support the hardware-assist takeover feature. Install a latest firmware for the hardware-assist takeover feature to work.

## Parameters

**hwtype** (STRING): Hardware module type.

## cf.hwassist.notifyCfgFailed

### Severity

ALERT

### Description

This message occurs when a controller module fails to update the hardware-assist configuration with its partner node. This is most likely because the SP component is not ready. The controller module will retry the hardware-assist configuration update automatically.

### Corrective Action

Use the `""system node service-processor show""` command to check whether the SP is online. The SP must be online to update the hardware-assist configuration. If the SP is operational and online, use the `""system node service-processor reboot-sp""` command to reboot the SP. If the message persists, contact NetApp technical support.

### Syslog Message

Failed to update the hardware-assist configuration with hardware component (%s): %s(%d).

### Parameters

**hwtype** (STRING): Hardware component type.

**errorstr** (STRING): Error string returned by the hardware component.

**error** (INT): Error number returned by the hardware component.

## cf.hwassist.notifyEnableOff

### Severity

INFORMATIONAL

### Description

This message occurs when the hwassist thread detects that 'notify enable' is turned off on the partner node.

### Corrective Action

None.

### Syslog Message

HA hw\_assist: hw\_assist functionality on the partner node has been disabled by the user.

**Parameters**

(None).

**cf.hwassist.notifyEnableOn****Severity**

INFORMATIONAL

**Description**

This message occurs when the user enables hw\_assist functionality on the partner node. hw\_assist will try to bind to the specified IP address and port.

**Corrective Action**

None.

**Syslog Message**

HA hw\_assist: hw\_assist functionality on the partner node has been enabled by the user.

**Parameters**

(None).

**cf.hwassist.ptnStartThrottle****Severity**

INFORMATIONAL

**Description**

This message occurs when partner has started throttling hw\_assist alerts. UDP socket will be closed and no alerts will be received.

**Corrective Action**

(None).

**Syslog Message**

Failover Monitor hw\_assist: Partner(%s) node has started throttling hw\_assist alerts.

**Parameters**

**partnerName** (STRING): Partner's name

**cf.hwassist.ptnStopThrottle****Severity**

INFORMATIONAL

**Description**

This message occurs when partner has stopped throttling hw\_assist alerts. hw\_assist will bind to UDP socket and wait for alerts.

**Corrective Action**

(None).

## Syslog Message

hw\_assist: Partner(%s) has stopped throttling hw\_assist alerts.

## Parameters

**partnerName** (STRING): Partner's name

## cf.hwassist.recvKeepAlive

### Severity

INFORMATIONAL

### Description

This message occurs when hwassist thread receives hw\_assist KeepAlive alert after missing at least two consecutive KeepAlive alerts.

### Corrective Action

(None).

## Syslog Message

hw\_assist: Received hw\_assist KeepAlive alert from partner(%s).

## Parameters

**partnerName** (STRING): Partner's name

## cf.hwassist.socBindFailed

### Severity

ERROR

### Description

This message occurs when hwassist thread fails to bind to a socket on a particular ipaddress. The error indicates why bind failed.

### Corrective Action

Check if specified port and IP address is valid. You may want to try with a different recommended port number.

## Syslog Message

hw\_assist: bind failed to port %d on IP address %s. Error %d

## Parameters

**port** (INT): Port on which bind was attempted.

**ipaddress** (STRING): IP address to which bind was attempted.

**error** (INT): error returned by bind call

## cf.hwassist.socInitFailed

### Severity

NOTICE

## Description

This message occurs when hwassist thread fails to create socket.

## Corrective Action

None.

## Syslog Message

hw\_assist: Fail to create socket.

## Parameters

(None).

## cf.hwassist.startThrottle

### Severity

NOTICE

### Description

This message occurs when the hardware assist subsystem thread detects that it is receiving hardware assist alerts at a rate higher than the maximum configurable rate and has temporarily stopped listening for hardware assist alerts. This can be an indication that another client is configured to send SNMP traps to the hardware assist trap IP address and port. If the partner node fails while listening for hardware assist alerts is disabled, detection of the partner failure can take slightly longer as loss of heartbeat from the partner or other non-hardware-assisted mechanism is instead needed to detect the failure. Subsequent receipt of the cf.hwassist.stopThrottle message indicates the hardware assist thread is again listening for hardware assist alerts.

### Corrective Action

Check whether some other client is configured to send SNMP alerts to the IP address and port on which hardware assist is configured. You can obtain the IP address using the ""storage failover hwassist show"" command. If this is the case, direct those SNMP alerts to the correct IP address and port. If not, contact NetApp technical support to further diagnose the issue.

### Syslog Message

hw\_assist: Received %llu hw\_assist alerts in %llu secs. Throttling alerts for %llu secs.

### Parameters

**num\_of\_traps** (LONGINT): Number of traps received.

**time\_interval** (LONGINT): Time window in which the traps were received.

**throttle\_time** (LONGINT): Time for which throttle will happen.

## cf.hwassist.stopThrottle

### Severity

NOTICE

### Description

This message occurs when a node has stopped throttling hardware assist alerts from the partner. This indicates the hardware assist thread has restarted listening for hardware assist alerts subsequent to the receipt of the cf.hwassist.startThrottle event.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor hw\_assist: Shutting off throttle for hw\_assist alerts from partner(%s).

### **Parameters**

**partnerName** (STRING): Partner's name.

## **cf.hwassist.takeoverTrapRecv**

### **Severity**

NOTICE

### **Description**

This message occurs when hwassist thread receives a hw\_assist alert which will result in partner takeover. EMS displays trap type as well as reason for the trap.

### **Corrective Action**

(None).

### **Syslog Message**

hw\_assist: Received takeover hw\_assist alert from partner(%s), %s because %s.

### **Parameters**

**partnerName** (STRING): Partner's name  
**trap\_type** (STRING): Type of trap received  
**trap\_event** (STRING): Type of trap event

## **cf.hwassist.unknownSig**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when hwassist thread receives an unknown signal.

### **Corrective Action**

None.

### **Syslog Message**

hw\_assist: Received unknown signal 0x%x.

### **Parameters**

**sigmask** (INT): Signal received by hwassist thread.

## **cf.ic events**



## **cf.ic.cc.config.changed**

### **Severity**

NOTICE

### **Description**

This message occurs when the controller configuration is changed from a single enclosure HA configuration to a dual enclosure HA configuration.

### **Corrective Action**

If the controller is part of a dual enclosure HA configuration, make sure that the psm-cc-config? environment variable is unset and reboot.

### **Syslog Message**

HA interconnect: Controller configuration was changed from single enclosure HA configuration to dual enclosure HA configuration.

### **Parameters**

(None).

## **cf.ic.clientinitFailed**

### **Severity**

ERROR

### **Description**

This message occurs when the interconnect client fails to initialize due to remote VI setup failure.

### **Corrective Action**

Reboot the controller. Contact NetApp technical support if the problem persists.

### **Syslog Message**

HA interconnect: Initialization failure for %s client.

### **Parameters**

**client** (STRING): Name of the client that failed to initialize.

## **cf.ic.disabled**

### **Severity**

ERROR

### **Description**

This message occurs when the HA interconnect is disabled by the ic-type environment variable.

### **Corrective Action**

Either remove the HA interconnect from the controller or change the ic-type environment variable.

### **Syslog Message**

HA interconnect: Card in slot %d has been disabled because of the ic-type environment variable.

**Parameters**

**slot** (INT): Slot that the interconnect is in.

**cf.ic.driver.flush.idWarn****Severity**

NOTICE

**Description**

This message occurs when the HA interconnect cannot process a flush request for a particular driver I/O request.

**Corrective Action**

(None).

**Syslog Message**

HA interconnect: Flush %s failed for channel %d and ID %d.

**Parameters**

**s** (STRING): Name of the error.

**nvidx** (INT): Interconnect channel.

**wrid** (INT): Work request ID.

**cf.ic.failed.memReg****Severity**

INFORMATIONAL

**Description**

This message occurs when memory registration fails.

**Corrective Action**

(None).

**Syslog Message**

HA interconnect: Memory registration operation failed at %d.

**Parameters**

**source** (INT): Source for the memory operation.

**cf.ic.flush.ACKWarn****Severity**

NOTICE

**Description**

This message occurs when the HA interconnect cannot send acknowledgments through its completion channel.

**Corrective Action**

(None).

**Syslog Message**

HA interconnect: Flush completion %s error.

**Parameters**

s (STRING): Name of the error.

**cf.ic.flush.MemError****Severity**

ERROR

**Description**

This message occurs during boot time initialization, when the HA interconnect encounters an error with memory registration or allocation.

**Corrective Action**

Reboot the controller to restart the HA interconnect. Contact NetApp technical support if the problem persists.

**Syslog Message**

HA interconnect: Flush memory %s error.

**Parameters**

s (STRING): Name of the error.

**cf.ic.flush.oodlinitFailed****Severity**

ERROR

**Description**

This message occurs when there is a failure in the initialization of a NIC to enable out-of-order delivery on the HA interconnect.

**Corrective Action**

Reboot the controller to enable HA interconnect flushing. Contact NetApp technical support if the problem persists.

**Syslog Message**

HA interconnect: Flush out-of-order delivery initialization failed.

**Parameters**

(None).

**cf.ic.heartBeatFailed**

**Severity**

ERROR

**Description**

This message occurs when heartbeats cannot be sent or received over the HA interconnect.

**Corrective Action**

Verify that both controllers in the HA pair are operational. If this message persists, disable and reenab the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands.

**Syslog Message**

HA interconnect: Heartbeat failed.

**Parameters**

(None).

**cf.ic.initFailed****Severity**

ERROR

**Description**

This message occurs when the driver fails to find an HA interconnect adapter or the interconnect is nonfunctional.

**Corrective Action**

Verify that the High Availability interconnect card is present and operational on the system. Reboot the controller. Contact NetApp technical support if this problem persists.

**Syslog Message**

HA interconnect: Unable to initialize process.

**Parameters**

(None).

**cf.ic.non.cc.config.changed****Severity**

NOTICE

**Description**

This message occurs when the controller configuration is changed from a dual enclosure HA configuration to a single enclosure HA configuration.

**Corrective Action**

If the controller is part of single enclosure HA configuration, make sure that the other controller is physically present in the chassis and that the psm-cc-config? environment variable is set to true, and then reboot.

## Syslog Message

HA interconnect: Controller configuration was changed from dual enclosure HA configuration to single enclosure HA configuration.

## Parameters

(None).

## cf.ic.notifyTimeout

### Severity

NOTICE

### Description

This message occurs when data transfer over the HA interconnect times out.

### Corrective Action

(None).

## Syslog Message

HA interconnect: %s transfer timed out.

## Parameters

**qp\_name** (STRING): Queue-pair name.

## cf.ic.portsCrossConnected

### Severity

ALERT

### Description

This message occurs when the HA interconnect driver determines that the ports between this controller and its partner are cross-connected.

### Corrective Action

Connect port 1 on this controller to port 1 on the partner, and connect port 2 on this controller to port 2 on the partner. If the interconnect cables have serial number tags on both ends, use the serial numbers to verify that the ports are properly connected. If the interconnect cables do not have serial number tags, disconnect one of the HA interconnect ports from one of the controllers. The LEDs of this port and one of the ports on the partner turn off. If the two ports have different port numbers, the cables are cross-connected. Reversing the connection on one of the controllers rectifies the problem.

## Syslog Message

HA interconnect: Port %d of this controller seems to be connected to port %d on the partner.

## Parameters

**localPort** (INT): Local port number.

**remotePort** (INT): Remote port number.

## cf.ic.probe.error

**Severity**

ALERT

**Description**

This event is generated when we fail to detect an High Availability (HA) interconnect communication device during system initialization. If this error occurs, internode communication within the HA pair is down.

**Corrective Action**

Examine the logs for messages regarding the HA interconnect. Resolve the HA interconnect problems.

**Syslog Message**

Failover monitor: no interconnect found WARNING: Failover monitor operating in degraded mode

**Parameters**

**errorCode** (INT): The internal error code

**cf.ic.sbb****Severity**

NOTICE

**Description**

This message occurs when a system SBB (Storage Bridge Bay) Compatibility event occurs. SBB Compatibility events determine the presence or absence of the HA partner controller on some platforms.

**Corrective Action**

(None).

**Syslog Message**

HA interconnect: SBB Compatibility Event. %s

**Parameters**

**event\_description** (STRING): Presence or absence of the partner controller and what action has been taken on the interconnect device.

**cf.ic.xferTimedOut****Severity**

ERROR

**Description**

This message occurs when data transfer over the HA interconnect times out. This is usually seen when the interconnect link goes down. It can also be encountered when the interconnect hardware, firmware, or driver experiences a hang.

**Corrective Action**

Contact NetApp technical support if the problem persists.

**Syslog Message**

HA interconnect: %s transfer timed out.

## Parameters

**qp\_name** (STRING): Name of the connection queue-pair that timed out.

## cf.ic.xferTimedOutVSA

### Severity

NOTICE

### Description

This message occurs when data transfer over the high-availability (HA) interconnect times out on virtual platforms. This is usually seen when the underlying network is busy or the system is operating under heavy load.

### Corrective Action

Contact NetApp technical support if the problem persists.

### Syslog Message

HA interconnect: %s transfer timed out.

## Parameters

**qp\_name** (STRING): Name of the connection queue-pair that timed out.

## cf.mccip events

### cf.mccip.med.auso.stDisabled

### Severity

ERROR

### Description

This message occurs when automatic switchover in a MetroCluster(tm) IP configuration is disabled because the mediator mailbox disk is not healthy.

### Corrective Action

Use the "storage failover show -fields local-mailbox-disks" command to verify that all nodes in the cluster can access their respective mailbox disks. Use the (privilege: advanced) "storage iscsi-initiator show" command to verify that iSCSI connections to the mediator are up and healthy.

### Syslog Message

Automatic switchover disabled: Mediator mailbox disk not healthy.

## Parameters

(None).

### cf.mccip.med.auso.stEnabled

### Severity

NOTICE

**Description**

This message occurs when automatic switchover in a MetroCluster(tm) IP configuration is enabled because the mediator mailbox disk is online.

**Corrective Action**

(None).

**Syslog Message**

Automatic switchover enabled.

**Parameters**

(None).

## cf.misc events

### cf.misc.givebackPartnerNotReady

**Severity**

NOTICE

**Description**

The message occurs when a giveback is initiated before the partner has reached waiting-for-giveback state.

**Corrective Action**

(None).

**Syslog Message**

Giveback is being initiated before the partner is in the waiting-for-giveback state. The partner is in the %s state.

**Parameters**

**state** (STRING): Current state of the partner.

### cf.misc.operatorDisable

**Severity**

ERROR

**Description**

This message occurs when the operator disables the high-availability (HA) pair.

**Corrective Action**

Use the "storage failover modify -enabled true" command to reenable failover.

**Syslog Message**

Failover monitor: operator initiated disabling of takeover.

**Parameters**

(None).



## **cf.misc.operatorDisableOption**

### **Severity**

ERROR

### **Description**

This event is issued when the operator disables a High Availability (HA) option.

### **Corrective Action**

Re-enable the option

### **Syslog Message**

(None).

### **Parameters**

**option** (STRING): Name of the option

## **cf.misc.operatorEnable**

### **Severity**

NOTICE

### **Description**

This event is issued when the operator enables failover between the 2 nodes.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: operator initiated enabling of failover.

### **Parameters**

(None).

## **cf.misc.operatorEnableOption**

### **Severity**

NOTICE

### **Description**

This event is issued when the operator enables a controller failover option.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**option** (STRING): Name of the option

## cf.misc.operatorForcedTakeover

### Severity

NOTICE

### Description

This event is issued when a forced takeover is initiated by the operator.

### Corrective Action

(None).

### Syslog Message

Failover monitor: forced takeover initiated by operator

### Parameters

(None).

## cf.misc.operatorGiveback

### Severity

INFORMATIONAL

### Description

This event is issued when a giveback is initiated by the operator.

### Corrective Action

(None).

### Syslog Message

Failover monitor: %s initiated by operator

### Parameters

**strengthString** (STRING): External string

**strengthCode** (INT): Internal code for strength

## cf.misc.operatorTakeover

### Severity

NOTICE

### Description

This event is issued when a takeover is initiated by the operator.

### Corrective Action

(None).

### Syslog Message

Failover monitor: takeover initiated by operator

**Parameters**

(None).

**cf.misc.ProgTakeover****Severity**

NOTICE

**Description**

This message is issued when a takeover is initiated programmatically by Data ONTAP®.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: takeover initiated.

**Parameters**

(None).

**cf.misc.ProgTakeoverFail****Severity**

NOTICE

**Description**

This message occurs when a takeover is initiated programmatically by Data ONTAP® and the takeover request fails.

**Corrective Action**

Check the reason string and repair the problem reported.

**Syslog Message**

Failover monitor: Programmatic takeover failed (%s)

**Parameters**

**errorstr** (STRING): Error string returned by the failover monitor.

**cf.misc.ProgTakeoverFailInit****Severity**

NOTICE

**Description**

This message occurs when a takeover is initiated programmatically by Data ONTAP® and the takeover request fails because the controller is set to "non\_ha" mode or is not initialized.

**Corrective Action**

Set the HA mode to "ha" by using the "storage failover modify -mode ha" command, and then reboot the node.

**Syslog Message**

Failover monitor: Programmatic takeover failed due to initialization issues.

**Parameters**

(None).

**cf.misc.ProgTakeoverFailInTO****Severity**

NOTICE

**Description**

This message occurs when a takeover is initiated programmatically by Data ONTAP® and the node is already in takeover mode.

**Corrective Action**

(None).

**Syslog Message**

Failover monitor: Programmatic takeover failed due to previous takeover.

**Parameters**

(None).

**cf.misc.ProgTakeoverFailNotUp****Severity**

NOTICE

**Description**

This message occurs when a negotiated takeover is initiated programmatically by Data ONTAP® and the partner node is not up.

**Corrective Action**

Bring up the partner node.

**Syslog Message**

Failover monitor: Programmatic takeover failed because the partner node is not up.

**Parameters**

(None).

**cf.misc.ProgTakeoverFailPartnerDeny****Severity**

NOTICE

**Description**

This message occurs when a takeover is initiated programmatically by Data ONTAP® and the partner node denies the takeover request.

**Corrective Action**

Check previous EMS messages to determine why takeover was denied.

**Syslog Message**

Failover monitor: Programmatic takeover failed because the partner node denied the takeover request.

**Parameters**

(None).

**cf.misc.ProgTakeoverFailShutdown****Severity**

NOTICE

**Description**

This message occurs when a takeover is initiated programmatically by Data ONTAP® and a shutdown is in progress.

**Corrective Action**

Verify that both nodes are booted.

**Syslog Message**

Failover monitor: Programmatic takeover failed because shutdown was in progress.

**Parameters**

(None).

**cf.mode events****cf.mode.auto.HA****Severity**

NOTICE

**Description**

This message occurs when the controller is automatically set to High Availability(HA) mode because of the configuration values in the Field Replaceable Unit(FRU) or the slot of the interconnect/NVRAM card (depending on the hardware type of the controller).

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

## **cf.mode.auto.non.HA**

### **Severity**

NOTICE

### **Description**

This message occurs when the controller is automatically set to non-High Availability (non-HA) mode because of the configuration values in the Field Replaceable Unit (FRU) or the slot of the interconnect/NVRAM card (depending on the hardware type of the controller).

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

(None).

## **cf.mode.HA**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the user sets the controller to High Availability(HA) mode through the command-line interface.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

(None).

## **cf.mode.HA.onboot**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the controller is set to High Availability (HA) mode according to options detected at boot.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

**Parameters**

(None).

**cf.mode.non.HA****Severity**

INFORMATIONAL

**Description**

This message occurs when the user sets the controller to non-High Availability (non-HA) mode through the command-line interface.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**cf.mode.non.HA.onboot****Severity**

INFORMATIONAL

**Description**

This message occurs when the controller is set to non-High Availability (non-HA) mode according to options detected at boot.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**cf.multidisk events****cf.multidisk.fatalProblem****Severity**

ERROR

**Description**

This message occurs when the node is waiting for the partner to take over, after the node encounters a multidisk error or some other fatal error.

**Corrective Action**

(None).

**Syslog Message**

Node encountered a multidisk error or other fatal error while waiting to be taken over. %s.

**Parameters**

**mdp\_msg** (STRING): Multi Disk Panic (MDP) string.

## cf.nducomplete events

**cf.nduComplete****Severity**

INFORMATIONAL

**Description**

This message occurs when a nondisruptive upgrade is completed successfully.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**token** (STRING): Unique token that identifies an instance of a nondisruptive upgrade.

**host** (STRING): Name of system that completed the nondisruptive upgrade.

## cf.nduincomplete events

**cf.nduIncomplete****Severity**

INFORMATIONAL

**Description**

This message occurs when a nondisruptive upgrade fails.

**Corrective Action**

Resolve the reason for the failure and retry the nondisruptive upgrade from the failed step.

**Syslog Message**

(None).

**Parameters**

**token** (STRING): Unique token that identifies an instance of a nondisruptive upgrade.

**host** (STRING): Name of the node that failed the nondisruptive upgrade.

**reason** (STRING): Reason for the upgrade failure.



# cf.ndustart events

## cf.nduStart

### Severity

INFORMATIONAL

### Description

This message occurs when a nondisruptive upgrade starts.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**token** (STRING): Unique token that identifies an instance of a nondisruptive upgrade.

**host** (STRING): Name of system that started the nondisruptive upgrade.

# cf.nm events

## cf.nm.nicReset

### Severity

NOTICE

### Description

This message occurs when the controller performs a soft reset of the HA interconnect card as part of automatic recovery from a timeout or link error.

### Corrective Action

(None).

### Syslog Message

HA interconnect: Initiating soft reset on card %d due to %s.

### Parameters

**nic\_num** (INT): Interconnect NIC number.

**reason** (STRING): Describes why the soft reset was initiated. Possible reasons: rendezvous jammed rendezvous fm timeout rendezvous peer-connect timeout

## cf.nm.nicViError

### Severity

INFORMATIONAL

### Description

This message occurs when any error is detected in any virtual interface while handling a link transition.

### Corrective Action

(None).

### Syslog Message

HA interconnect: NIC %d has an error on %s (virtual interface #%d): %s.

### Parameters

**nic\_num** (INT): Interconnect NIC number.  
**vi\_name** (STRING): Name of the virtual interface.  
**vinum** (INT): Virtual interface number.  
**origin** (STRING): Origin of the error.

## cf.noautogb events

### cf.noAutoGB.ptnrLowerVersion

#### Severity

NOTICE

#### Description

This message occurs when an automatic giveback operation is cancelled because the high-availability (HA) partner is using an earlier and incompatible Data ONTAP® version than the local node. Subsequent failover events might not be successful when HA partners are not using compatible Data ONTAP versions.

#### Corrective Action

Run the "system node image show" command to determine which image on the partner should be booted to be compatible with the local node. After rebooting the node, automatic giveback is performed.

#### Syslog Message

Auto-giveback cancelled because the HA partner is using a lower Data ONTAP version than the local node.

#### Parameters

(None).

## cf.nopartner\_takeover events

### cf.noPartner\_takeover

#### Severity

ERROR

#### Description

This message occurs when the partner node fails but requested not to be taken over to prevent multiple, simultaneous RAID reconstructs from starting.

#### Corrective Action

Check the partner node for multiple disk or disk shelf failures, and then correct that situation. If you want takeover to occur, use the "storage failover takeover" command. If you want takeover to occur when disk or shelf failures occur, change the setting of the "raid.panic.missing.disks" option to '0' from the nodeshell.

### **Syslog Message**

Failover monitor: Partner node is down, but requested not to be taken over.

### **Parameters**

(None).

## **cf.nullshelfcountmsg events**

### **cf.nullShelfCountMsg**

#### **Severity**

ERROR

#### **Description**

This event occurs if we are sent a NULL negotiated failover shelf count message. \* This is likely due to an InterConnect problem.

#### **Corrective Action**

Check for and resolve any interconnect errors.

### **Syslog Message**

Invalid disk shelf count message was received from partner.

### **Parameters**

(None).

## **cf.partner events**

### **cf.partner.login**

#### **Severity**

INFORMATIONAL

#### **Description**

This event is issued when a login to the console of the partner occurs as a result of the 'partner' command being issued with no arguments.

#### **Corrective Action**

(None).

### **Syslog Message**

Login to partner shell: %s

### **Parameters**

**hostname** (STRING): The name of this node

### **cf.partner.logoff**

**Severity**

INFORMATIONAL

**Description**

This event is issued when a logoff from the console of the partner occurs as a result of the 'partner' command being issued with no arguments.

**Corrective Action**

(None).

**Syslog Message**

Logoff from partner shell: %s

**Parameters**

**hostname** (STRING): The name of this node

**cf.partner.nvram.notSync****Severity**

INFORMATIONAL

**Description**

This message occurs during takeover, when the system finds the partner nonvolatile memory (NVRAM) unsynchronized. There might be data loss.

**Corrective Action**

(None).

**Syslog Message**

Partner NVRAM was not synchronized. Some data might be lost.

**Parameters**

(None).

**cf.partner.nvram.state****Severity**

INFORMATIONAL

**Description**

This message occurs during takeover, when the system determines that the partner mailbox is stale. This might mean that partner nonvolatile memory (NVRAM) is unsynchronized and might have some data loss.

**Corrective Action**

(None).

**Syslog Message**

Partner mailbox was stale. Partner NVRAM might not be synchronized and some data might be lost.

**Parameters**

(None).

## cf.partner.ready.giveback

### Severity

INFORMATIONAL

### Description

This message occurs when the partner is booted and ready for giveback.

### Corrective Action

Verify that network connectivity is restored, and then use the "storage failover giveback" command to initiate giveback to the partner.

### Syslog Message

Partner is booted and ready for giveback.

### Parameters

(None).

## cf.partner.shortUptime

### Severity

ERROR

### Description

This message occurs when a node in a controller failover pair determines that the partner node has stayed up for a very short time period. This might result in a takeover

### Corrective Action

Check the disk connectivity to the partner node. If the problem persists and the partner node restarts multiple times, perform hardware diagnostics in maintenance mode, and then consult the EMS logs to determine why this is happening.

### Syslog Message

Partner up for %llu seconds only.

### Parameters

**seconds** (LONGINT): Time duration for which the partner was up.

## cf.reservation events

### cf.reservation.disk.notFound

### Severity

NOTICE

### Description

This message occurs when the system does not halt on a disk reservation conflict and the reservation is not found on the disk.

**Corrective Action**

(None).

**Syslog Message**

Reservation found on disk %s; not halting due to sanown\_pr\_test\_mode.

**Parameters**

**diskName** (STRING): Disk name.

## cf.rsrc events

### cf.rsrc.givebackFail

**Severity**

ALERT

**Description**

This message occurs when a failure is detected during giveback. The giveback operation will be cancelled.

**Corrective Action**

Check the logs for other messages regarding the listed subsystem using the "event log show" CLI command. Resolve those problems and perform a planned giveback operation using "storage failover giveback" command.

**Syslog Message**

Failover monitor: giveback during %s failed; giveback cancelled.

**Parameters**

**op** (STRING): Name of the giveback subsystem that has suffered a fatal error.

### cf.rsrc.givebackForceFail

**Severity**

ERROR

**Description**

This event occurs if a failure is forced during giveback. This event should only occur in debugging kernels.

**Corrective Action**

Retry the takeover.

**Syslog Message**

Failover monitor: giveback during %s forced fail

**Parameters**

**op** (STRING): Name of the takeover operation

### cf.rsrc.givebackForceVeto

**Severity**

ERROR

**Description**

This event occurs if a veto is forced during giveback. This event should only occur in debugging kernels.

**Corrective Action**

Check the state of the subsystem listed in this message. Retry the giveback.

**Syslog Message**

Giveback during %s forced veto

**Parameters**

**op** (STRING): Name of the takeover operation

**cf.rsrc.givebackOpFail****Severity**

ERROR

**Description**

This event occurs if a failure is detected during giveback. This error is not fatal to the giveback process but indicates that the related service could not be shutdown cleanly.

**Corrective Action**

Check the state of the subsystem listed in the message and validate it is working correctly. Validate that there are no configuration errors for the listed subsystem.

**Syslog Message**

Failover monitor: giveback during %s failed; giveback continuing...

**Parameters**

**op** (STRING): Name of the giveback operation that has suffered a non-fatal error.

**cf.rsrc.givebackUnexpected****Severity**

ALERT

**Description**

This event occurs if we detect an unexpected failover monitor state during a giveback operation. This indicates a software bug.

**Corrective Action**

Retry the giveback.

**Syslog Message**

(None).

**Parameters**

**op** (STRING): Name of the giveback operation

**state** (INT): Value of the failover monitor state

## **cf.rsrc.givebackVeto**

### **Severity**

ALERT

### **Description**

This message occurs when a veto is indicated during giveback. A veto occurs when the active partner state is detected during the giveback process. A veto can be overridden by using the "cf giveback -f" or "storage failover takeover -override-vetoes true" command.

### **Corrective Action**

Check the preceding messages for reasons why the giveback was vetoed. Check the state of the subsystem listed in this message. Retry the giveback.

### **Syslog Message**

Failover monitor: %s: giveback canceled due to active state.

### **Parameters**

**op** (STRING): Name of the takeover operation.

## **cf.rsrc.raidReplayOnlyTODone**

### **Severity**

INFORMATIONAL

### **Description**

This event occurs when a node is booting after having panicked too soon after takeover or during takeover. In that event we want to replay the partner RAID logs, before we perform a giveback. This message indicates the point at which we have committed partner RAID logs to the disk and are now going to initiate a giveback.

### **Corrective Action**

(None).

### **Syslog Message**

Failover monitor: Raid replay-only takeover completed. Initiating giveback.

### **Parameters**

(None).

## **cf.rsrc.replayOnlyTakeoverDone**

### **Severity**

NOTICE

### **Description**

This event occurs when a node is booting after having panicked too soon after takeover or during takeover. In that event we want to replay the partner logs, before we perform a giveback. This message indicates the point at which we have committed partner logs to the disk and are now going to initiate a giveback.



**Corrective Action**

(None).

**Syslog Message**

Failover monitor: replay-only takeover completed. Initiating giveback.

**Parameters**

(None).

**cf.rsrc.takeoverFail****Severity**

ALERT

**Description**

This message occurs when a failure is detected during takeover. The takeover operation will be cancelled.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Failover monitor: takeover during %s failed; takeover cancelled.

**Parameters**

**op** (STRING): Name of the takeover operation that has suffered a fatal error.

**cf.rsrc.takeoverForceFail****Severity**

ERROR

**Description**

This event occurs if a failure is forced during takeover. This event should only occur in debugging kernels.

**Corrective Action**

Retry the takeover request.

**Syslog Message**

Failover monitor: takeover during %s forced fail

**Parameters**

**op** (STRING): Name of the takeover operation

**cf.rsrc.takeoverOpFail****Severity**

ERROR

**Description**

This event occurs if a failure is detected during takeover. This error is not fatal to the takeover process but indicates that the related service may be inaccessible during takeover of the partner.

### **Corrective Action**

After takeover completes, check the state of the subsystem listed in the message. Check that the configuration of the subsystem listed in the message is correct.

### **Syslog Message**

Failover monitor: takeover during %s failed; takeover continuing...

### **Parameters**

**op** (STRING): Name of the takeover operation that has suffered a non-fatal error.

## **cf.rsrc.takeoverUnexpected**

### **Severity**

ERROR

### **Description**

This event occurs if we detect an unexpected failover monitor state during a takeover operation. This indicates a software bug.

### **Corrective Action**

Check the state of the subsystem listed in 'op' parameter. Retry the giveback.

### **Syslog Message**

(None).

### **Parameters**

**op** (STRING): Name of the giveback operation

**state** (INT): Value of the failover monitor state

## **cf.rv events**

### **cf.rv.connected**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an HA connection agent connects on an HA interconnect device.

### **Corrective Action**

(None).

### **Syslog Message**

HA interconnect: Connection for '%s' succeeded.

### **Parameters**

**rv\_name** (STRING): Logical name of the successful HA connection.

## cf.rv.flush.handleExchange

### Severity

INFORMATIONAL

### Description

This message occurs when HA interconnect exchanges flushing information with the partner.

### Corrective Action

If the flush handle exchange reports "inactive", the remote controller might be running a version of Data ONTAP® that does not support flushing. Make sure that the controller and its partner are running compatible versions of the software.

### Syslog Message

HA interconnect: Flushing is %s.

### Parameters

s (STRING): Flushing status.

## cf.rv.localNoConn

### Severity

INFORMATIONAL

### Description

This message occurs when the local controller votes not to connect to the HA partner. This can be caused by incompatibility between the software versions running on the two controllers.

### Corrective Action

Install compatible versions of software on both controllers.

### Syslog Message

HA interconnect: Local controller voted not to connect.

### Parameters

(None).

## cf.rv.nicReset

### Severity

ERROR

### Description

This message occurs when HA interconnects are reset due to NVRAM log being unsynchronized for longer time than 15seconds.

### Corrective Action

Contact NetApp technical support if the problem persists.

### Syslog Message

HA interconnect: Reset due to unsynchronized log.

## Parameters

(None).

## cf.rv.notConnected

### Severity

ALERT

### Description

This message occurs when an HA connection agent fails to connect on an HA interconnect device.

### Corrective Action

Verify that both controllers in the HA pair are operational. If this message persists, disable and reenable the link by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands.

### Syslog Message

HA interconnect: Connection for '%s' failed.

### Parameters

**rv\_name** (STRING): Logical name of the failing HA connection.

## cf.rv.partnerNoConn

### Severity

ERROR

### Description

This message occurs when the HA partner votes not to connect to this controller. This can happen if the HA partner detects an incompatibility between the software versions running on the two controllers.

### Corrective Action

Install compatible versions of software on both controllers.

### Syslog Message

HA interconnect: Partner voted not to connect.

### Parameters

(None).

## cf.shutdown events

### cf.shutdown.info

### Severity

NOTICE

### Description

This message occurs when the controller receives a message from its partner over the HA interconnect that indicates a condition related to the partner's attempted graceful shutdown.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**eventType** (STRING): Event type, such as "Long shutdown".

**eventSummary** (STRING): Last task executed and its duration.

**shutdownSummary** (STRING): Shutdown phase summary for the event.

## cf.sysid events

### cf.sysid.sent

#### Severity

INFORMATIONAL

#### Description

This message occurs when a node sends its system ID and its HA partner's system ID successfully across the HA interconnect.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**sysid** (LONGINT): Value of the local system ID that was sent successfully .

**host** (STRING): This parameter is obsolete from Data ONTAP®8.3 and later because the partner's systemID is also sent in the message.

**partner\_sysid** (LONGINT): Value of the partner system ID that was sent successfully.

## cf.takeover events

### cf.takeover.aborted.bdfu

#### Severity

NOTICE

#### Description

This message occurs when negotiated takeover is aborted because the background disk firmware update took more than 120 seconds to finish.

#### Corrective Action

Rerun the takeover command.

**Syslog Message**

(None).

**Parameters**

(None).

**cf.takeover.delayed.bdfu****Severity**

INFORMATIONAL

**Description**

This message occurs when negotiated takeover is initiated on a node while the background disk firmware update (BDFU) is in progress. If update is not completed within 120 seconds, takeover is aborted.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**cf.takeover.delayed.cleanup****Severity**

INFORMATIONAL

**Description**

This message occurs when takeover of a node by the node's high-availability (HA) partner was delayed due to the cleanup being performed after the previous giveback event.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**delay** (LONGINT): Amount of time (in milliseconds) the takeover was delayed.

**cf.takeover.disabled****Severity**

ALERT

**Description**

This message occurs when takeover of a node by the node's high-availability (HA) partner is disabled.

## Corrective Action

Resolve the issue, as identified by the message, for takeover being disabled. Use the "storage failover show -instance" command to determine whether the disks are accessible by both nodes in the HA pair, also check the HA interconnect between both nodes is installed and functioning correctly. If not, fix the cabling, disk, or adapter issues to resolve the connectivity issues.

## Syslog Message

%s.

## Parameters

**reason** (STRING): Reason for takeover being disabled.

## cf.takeover.missing.ptnrDiskInventory

### Severity

EMERGENCY

### Description

This message occurs when a takeover request is denied because the node did not receive the partner disk inventory information.

## Corrective Action

Wait five minutes and then try takeover again. If you want to disable the disk inventory check instead, you can use the "cf disable diskInventory" command from the advanced privilege level of the nodeshell.

However, disabling the disk inventory check might lead to a client outage when the local node cannot see the partner disks during takeover.

## Syslog Message

(None).

## Parameters

(None).

## cf.takeover.missingPtnrDisks

### Severity

EMERGENCY

### Description

This message occurs when a takeover request is denied because the node cannot see some of the partner's drives.

## Corrective Action

Make sure that the local node has connectivity to all file system drives owned by the partner. Use the "storage failover show -fields local-missing-disks, partner-missing-disks" command to rescan the drive and to determine latest ownership. Drives might not be visible if partner drives have failed or the node is having issues accessing the partner node's drives. Verify that the HA interconnect cabling is correct, replace any failed drives and then check whether the issue is resolved. If you want to disable the drive inventory check instead, use the "storage failover takeover -allow-disk-inventory-mismatch true" command. However, disabling the drive inventory check during takeover is not advisable, because it might lead to client outage.

## Syslog Message

Failover monitor: Takeover failed because the node cannot see some of the partner node's drives.

## Parameters

(None).

# cf.transition events

## cf.transition.info

### Severity

NOTICE

### Description

This message occurs when a failover event is completed successfully. It displays the protocol transition information for the event.

### Corrective Action

(None).

## Syslog Message

%s %s.

## Parameters

**eventType** (STRING): Takeover, aggregate relocation, SFO phase of takeover, or SFO giveback event type.

**eventSummary** (STRING): Transition duration for protocols during the event.



# cifs events

## cifs.does events

### cifs.does.not.exist

#### Severity

ERROR

#### Description

This message occurs when a Vserver or Cluster administrator tries to initialize Snapmirror on Async VserverDR ID Discard setup without creating a CIFS server at destination. Snapmirror initialize will not complete successfully and will remain uninitialized.

#### Corrective Action

Create CIFS server at destination(DP-destination) vservers. Reissue snapmirror initialize.

#### Syslog Message

CIFS server does not exist on Vserver '%s'.

#### Parameters

**vserverName** (STRING): Name of the DP-destination Vserver on which CIFS server does not exist.

## cifs.domainpwd events

### cifs.domainpwd.job.deletion

#### Severity

NOTICE

#### Description

This message occurs when disabling of CIFS domain password change schedule fails to delete the job for a Vserver.

#### Corrective Action

Delete the job on the owning node using 'job delete -id ' command.

#### Syslog Message

An attempt to delete the CIFS domain password change job with ID %d for Vserver %s failed. Reason %s.

#### Parameters

**jobid** (LONGINT): CIFS password change job ID that failed to delete.

**vserver** (STRING): Name of the Vserver for which the password change job failed to delete.

**reason** (STRING): Reason for the failure to delete the job.

### cifs.domainpwd.job.not.sched

#### Severity

ERROR

## Description

This message occurs when the scheduled job for a CIFS domain password change fails for a Vserver.

## Corrective Action

Disable and enable the feature for the Vserver by using the 'vserver cifs domain password schedule modify -is-schedule-enabled false' and 'vserver cifs domain password schedule modify -is-schedule-enabled true' commands. If the error persists, do the following: 1. Disable the scheduling feature for the password change on the Vserver by using the command specified above. 2. Change password manually by using the 'vserver cifs domain password change' command. 3. Contact NetApp technical support.

## Syslog Message

An attempt to schedule a job to update the domain account password for Vserver %s failed.

## Parameters

**vserver** (STRING): Name of the Vserver for which the password change job failed.

## cifs.domainpwd.not.updated

### Severity

ERROR

### Description

This message occurs when the CIFS domain password update for a Vserver fails.

### Corrective Action

Try updating the password manually by using the 'vserver cifs domain password change' command. Review recent event logs for additional information. If the error persists, contact NetApp technical support.

### Syslog Message

An attempt to update the domain account password for Vserver %s failed during %s with the following error:  
%s

### Parameters

**vserver** (STRING): Name of the Vserver for which the password update failed.  
**method** (STRING): Method used to update the password.  
**reason** (STRING): Reason for the failure.

## cifs.lug events

### cifs.lug.import.fail

#### Severity

ERROR

#### Description

This message occurs when the job to import CIFS local users and groups is completed successfully only at a primary site within a MetroCluster® configuration or Vserver DR setup. CIFS local users and groups have either not been applied or only partially applied at the secondary site.

### **Corrective Action**

Use the "metrocluster vserver resync -vserver <vserver name>" command to resynchronize the Vserver configuration.

### **Syslog Message**

CIFS local users and groups import job failed for Vserver %s. Reason: %s.

### **Parameters**

**vserver** (STRING): Name of the Vserver.

**reason** (STRING): Reason for the failure.

## **cifs.replay events**

### **cifs.replay.failure**

#### **Severity**

ALERT

#### **Description**

This message occurs when configuration is not synchronized. This might cause the node experiencing the problem to become unresponsive to data access.

#### **Corrective Action**

Contact NetApp technical support. Be prepared to provide log files to technical support, particularly the "mgwd.log" and "sktrace.log" files.

#### **Syslog Message**

Failed to replay configuration for %s module.

#### **Parameters**

**module** (STRING): Name of the module for which configuration is not synchronized.

## **cifs.server events**

### **cifs.server.auth.StyleMismatch**

#### **Severity**

ERROR

#### **Description**

This message occurs when the Vserver DR subsystem detects that the authentication style (domain or workgroup) of the source Vserver's CIFS server does not match the authentication style of the destination Vserver's CIFS server

#### **Corrective Action**

Use the 'vserver cifs modify' command to configure the source or destination Vserver's CIFS server into a domain or workgroup as necessary for the authentication styles to match.

## Syslog Message

Vserver DR check found that the CIFS server authentication style of the source Vserver, %s, did not match the authentication style of the destination Vserver, %s.

## Parameters

**source\_vserver** (STRING): Name of the source Vserver of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

# cifs.shadowcopy events

## cifs.shadowcopy.abort

### Severity

NOTICE

### Description

This message occurs when a shadow copy set is automatically aborted when no activity has been detected for a certain period of time. This condition could be caused by improper operation of the backup application and/or the Volume Shadow Copy Service (VSS) client, or the loss of network connection to the host.

### Corrective Action

Check the backup application and VSS client. Verify the network connection between the file server and the client.

## Syslog Message

The shadow copy set ( Client Shadow Copy Set ID : %s , Filer Shadow Copy Set ID : %s ) has been automatically aborted. No activity has been seen for at least %d minutes.

## Parameters

**clientShadowCopySetId** (STRING): Client Shadow Copy Set ID.

**filerShadowCopySetId** (STRING): File Server Shadow Copy Set ID.

**timeout** (INT): No activity timeout in minutes.

## cifs.shadowcopy.failure

### Severity

ERROR

### Description

This message occurs when a shadow copy operation has failed.

### Corrective Action

Using the information provided in the event message, here are some things to check on your system: Shadow copy configuration; Is the feature enabled? Are the appropriate licenses installed? Share(s) on which the shadow copy operation is performed; Is the share name correct? Does the share path exist? State of the shadow copy set and its shadow copies.

## Syslog Message

A shadow copy operation has failed: %s. ( Operation : %s , Client Shadow Copy Set ID : %s , Filer Shadow Copy Set ID : %s , Client Shadow Copy ID : %s , Filer Shadow Copy ID : %s , Share Name : %s, Object type is: %s with UUID: %s )

## Parameters

**errMsg** (STRING): Reason the Shadow Copy operation failed.  
**operation** (STRING): Shadow Copy operation that failed.  
**clientShadowCopySetId** (STRING): Client Shadow Copy Set ID.  
**filerShadowCopySetId** (STRING): File Server Shadow Copy Set ID.  
**clientShadowCopyId** (STRING): Client Shadow Copy ID.  
**filerShadowCopyId** (STRING): File Server Shadow Copy ID.  
**shareName** (STRING): Share Name.  
**object\_type** (STRING): Type of resource object.  
**object\_uuid** (STRING): UUID of the resource object.

# cifs.static events

## cifs.static.homedir.create

### Severity

ERROR

### Description

This message occurs when a user creates a static home directory without a substitution value (%u or %w) to uniquely identify a home directory share. If multiple users from the same client access their home directories at the same time, unprotected access to each other's files can occur. This might cause unexpected \"access denied\" errors on those files, and data corruption within those files. Note: multiple users from a single client is most likely in a Citrix or Windows Terminal Services environment.

### Corrective Action

Ensure that a substitution value (%u or %w) is provided to the name of the home directory share.

### Syslog Message

A static home directory share \"%s\" without a substitution value (%%u or %%w) is created on Vserver: %s.

### Parameters

**shareName** (STRING): Name of the share that the user has created.  
**vserverName** (STRING): Name of the Vserver associated with this operation.

# clam events

## clam.disable events

### clam.disable

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) is disabled by either turning off the KCS capability or by setting the 'clam sysctl/bootarg' variable.

#### Corrective Action

If you want CLAM functionality disabled, no action is needed. Otherwise, reenable CLAM and reboot the node to make CLAM operational again.

#### Syslog Message

CLAM functionality is disabled.

#### Parameters

(None).

### clam.disable.takeover

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) is disabled from initiating storage failover requests. This capability is automatically disabled in a NAS-only configuration. To control the setting of CLAM takeover, additional interfaces, such as sysctl bootarg, are provided.

#### Corrective Action

No action is needed.

#### Syslog Message

CLAM is disabled from performing storage failover.

#### Parameters

(None).

## clam.enable events

### clam.enable

#### Severity

INFORMATIONAL

## Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) is enabled to operate on this node. When all nodes in the cluster are upgraded to a KCS-available image, KCS capability gets turned on. This event is an indication that CLAM will be now begin to function.

## Corrective Action

No action is needed.

## Syslog Message

CLAM functionality is enabled.

## Parameters

(None).

## clam.enable.takeover

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) is enabled to initiate storage failover requests when the CLAM OOQ (Out Of Quorum) condition is detected on the partner node. This capability is automatically enabled in a scalable SAN configuration. To control the setting of CLAM takeover, additional interfaces, such as `sysctl bootarg`, are provided.

### Corrective Action

No action is needed.

### Syslog Message

CLAM is enabled to perform storage failover when partner node is OOQ.

### Parameters

(None).

## clam.failover events

### clam.failover.event

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) receives a failover event from high-availability(HA) subsystem.

#### Corrective Action

(None).

#### Syslog Message

CLAM on node (id=%llu) received an HA failover event for partner (ID=%llu) indicating that %s.

## Parameters

**localnodeid** (LONGINT): Node ID of the local node.

**partnernodeid** (LONGINT): Node ID of the partner node.

**event\_description** (STRING): Description of the event that CLAM received from the HA subsystem.

# clam.heartbeat events

## clam.heartbeat.state.change

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services Connectivity, Liveliness, Availability Monitor (CLAM) detects that the state of the heartbeats to a node changed(that is, either the heartbeats started succeeding or they started failing).

### Corrective Action

Take one or more of the following actions: - A repeated toggling of the heartbeat state indicates intermittent network issues. Check the connectivity of the node to the cluster interconnect and fix any problems. - In all the other cases, the heartbeat failure will eventually result in the node going out of quorum. If that happens, take the actions prescribed for "EMS\_clam\_node\_oq event".

### Syslog Message

Heartbeats to node (name=%s, ID=%llu) are %s.

## Parameters

**hostname** (STRING): Host name of the node.

**nodeid** (LONGINT): Node ID of the node.

**state** (STRING): New state of the heartbeats to this node.

# clam.invalid events

## clam.invalid.config

### Severity

ERROR

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) determines that the local node is in an invalid configuration for providing CLAM functionality.

### Corrective Action

Take one or more of the following actions, as appropriate: - Enable failover on the node by running the 'storage failover modify -node node\_name -enabled true' command. - If there is a version incompatibility between the local node and the other nodes in the cluster, upgrade the nodes as needed.

### Syslog Message

Local node (name=%s, id=%llu) is in an invalid configuration for providing CLAM functionality. CLAM cannot determine the identity of the HA partner.



## Parameters

**localhostname** (STRING): Host name of the local node.

**localnodeid** (LONGINT): Node ID of the local node.

# clam.mailbox events

## clam.mailbox.initialized

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) has written to the high availability (HA) mailboxes that it has initialized and is ready to participate in CLAM quorum.

### Corrective Action

(None).

### Syslog Message

CLAM on local node has initialized its mailboxes and is ready to participate in quorum.

## Parameters

(None).

# clam.node events

## clam.node.disabled

### Severity

INFORMATIONAL

### Description

This message occurs when the local Connectivity, Liveliness and Availability Monitor (CLAM) concludes that the local node is now disabled.

### Corrective Action

This EMS is logged when the local node has lost connectivity with the rest of the nodes in the cluster and the connectivity has not been restored in time for the local node to participate in cluster quorum. Take the actions prescribed for "EMS\_clam\_node\_ooq" event.

### Syslog Message

CLAM published an event indicating the local node is now disabled(reason: %s).

## Parameters

**disabled\_reason** (STRING): Reason for the node getting disabled.

## clam.node.inq

**Severity**

INFORMATIONAL

**Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) determines that a node in the cluster has transitioned to "cluster quorum".

**Corrective Action**

(None).

**Syslog Message**

Cluster node (name=%s, ID=%llu) is in "CLAM quorum".

**Parameters**

**hostname** (STRING): Host name of the cluster node that transitioned to quorum.

**nodeid** (LONGINT): Node ID of the cluster node that transitioned to quorum.

**clam.node.oog****Severity**

EMERGENCY

**Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) determines that a node in the cluster is out of "cluster quorum".

**Corrective Action**

Take one or more of the following actions: - Follow the KB article 2022109 to troubleshoot the resolve the issue. - Check the connectivity of the node to the cluster interconnect and fix any problems. - If the local node is administratively disabled, enable it by running the following command from the clustershell 'cluster modify -node node\_name -eligibility true'. - If the node is supposed to be in an HA pair but is not, enable the HA pair. You can check whether a node is in an HA pair by running 'storage failover show' command from the clustershell. You can enable storage failover by running the 'storage failover modify -enabled true' command from the clustershell. - If there is a version incompatibility between the local node and the other nodes in the cluster, upgrade the nodes as needed. You can see the software version of all the nodes from the clustershell using the 'version -node \*' command.

**Syslog Message**

Node (name=%s, ID=%llu) is out of "CLAM quorum" (reason=%s).

**Parameters**

**hostname** (STRING): Host name of the node that is out of quorum.

**nodeid** (LONGINT): Node ID of the node that is out of quorum.

**reason** (STRING): Reason for the out of quorum transition.

**clam.node.paused****Severity**

INFORMATIONAL

**Description**

This message occurs when the local Connectivity, Liveliness and Availability Monitor (CLAM) entered a paused state.

**Corrective Action**

(None).

**Syslog Message**

CLAM published an event indicating the local node has paused.

**Parameters**

(None).

**clam.node.resumed****Severity**

INFORMATIONAL

**Description**

This message occurs when the local Connectivity, Liveliness and Availability Monitor (CLAM) was previously in a paused state and has now resumed.

**Corrective Action**

(None).

**Syslog Message**

CLAM published an event indicating the local node has resumed.

**Parameters**

(None).

**clam.node.start.voting****Severity**

INFORMATIONAL

**Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) on a local node has started to vote for its high-availability (HA) partner node.

**Corrective Action**

(None).

**Syslog Message**

CLAM on local node (ID:%llu) has started to vote for its HA partner node (ID:%llu).

**Parameters**

**local\_id** (LONGINT): Node ID of the local node.

**partner\_id** (LONGINT): Node ID of the partner node.

## clam.node.stop.voting

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) on a local node has stopped voting for its high-availability (HA) partner node.

### Corrective Action

(None).

### Syslog Message

CLAM on local node (ID:%llu) has stopped to vote for its HA partner node (ID:%llu).

### Parameters

**local\_id** (LONGINT): Node ID of the local node.

**partner\_id** (LONGINT): Node ID of the partner node.

## clam.partner events

### clam.partner.halting

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveness, Availability Monitor (CLAM) is notified that its high-availability (HA) partner node is undergoing a planned shutdown.

#### Corrective Action

(None).

#### Syslog Message

CLAM is notified that its HA partner node is undergoing a planned shutdown (reason: %s).

#### Parameters

**reason** (STRING): Reason for the partner beginning the halt process.

## clam.peer events

### clam.peer.halting

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveness, and Availability Monitor (CLAM) is notified that its high-availability (HA) partner node is undergoing a planned shutdown.

### **Corrective Action**

(None).

### **Syslog Message**

CLAM is notified that node %s (ID=%llu) is undergoing a planned shutdown.

### **Parameters**

**hostnode** (STRING): Name of node beginning the halt process.

**nodeid** (LONGINT): ID of node beginning the halt process.

## **clam.reenable events**

### **clam.reenable**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) is reenabled to operate on this node after being intentionally disabled. The intentional disabling could have been through turning off KCS capability or setting the 'clam bootarg/sysctl' variable. Reenabling CLAM functionality warrants a reboot, upon which CLAM will be functional.

#### **Corrective Action**

Reboot the node to make CLAM operational.

#### **Syslog Message**

CLAM functionality is reenabled. Reboot the node to make CLAM operational.

#### **Parameters**

(None).

## **clam.start events**

### **clam.start.mailbox.polling**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) has started a poll to read CLAM mailboxes.

#### **Corrective Action**

(None).

#### **Syslog Message**

CLAM on local node (ID=%llu) has started a poll to read HA partner (ID=%llu) mailbox.

## Parameters

**local\_id** (LONGINT): Node ID of the local node.  
**partner\_id** (LONGINT): Node ID of the partner node.

# clam.stop events

## clam.stop.mailbox.polling

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) has stopped polling the CLAM mailboxes.

### Corrective Action

(None).

### Syslog Message

CLAM on local node (ID=%llu) has stopped the poll to read HA partner (ID=%llu) mailbox.

## Parameters

**local\_id** (LONGINT): Node ID of the local node.  
**partner\_id** (LONGINT): Node ID of the partner node.

# clam.takeover events

## clam.takeover

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services Connectivity, Liveliness, Availability Monitor (CLAM) triggers takeover of a partner node after determining that the local node is in "CLAM quorum" and the partner node is out of "CLAM quorum".

### Corrective Action

Take one or more of the following actions: - If the node panicked or is hung, troubleshoot and resolve the problem. - Check the connectivity of the node to the cluster interconnect and fix any problems. - If the local node is administratively disabled, enable it by running the following command from the clustershell 'cluster modify -node node\_name -eligibility true'. - If the node is supposed to be in an HA pair but is not, enable the HA pair. You can check whether a node is in an HA pair by running the 'storage failover show' command from the clustershell. You can enable storage failover by running the 'storage failover modify -enabled true' command from the clustershell. - If there is a version incompatibility between the local node and the other nodes in the cluster, upgrade the nodes as needed. You can see the software version of all the nodes from the clustershell using the 'version -node \*' command.

### Syslog Message

Local node (name=%s, ID=%llu) initiated takeover of partner node (name=%s, ID=%llu, state=%llu)

result=%s, type=%s.

## Parameters

**localhostname** (STRING): Host name of the local node.  
**localnodeid** (LONGINT): Node ID of the local node.  
**partnerhostname** (STRING): Host name of the partner node.  
**partnernodeid** (LONGINT): Node ID of the partner node.  
**partnerstate** (LONGINTEX): The high-availability state of the partner node  
**result** (STRING): Result of the takeover initiation operation.  
**type** (STRING): Takeover Type.

## clam.takeover.disallowed

### Severity

INFORMATIONAL

### Description

This message occurs when Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) cannot initiate takeover operations. Takeover is disallowed by High Availability (HA) subsystem on the partner node.

### Corrective Action

(None).

### Syslog Message

CLAM on node %s (ID=%llu) cannot proceed with a takeover operation of partner node %s (ID=%llu) because HA disallowed takeover. Reason :%s

## Parameters

**local\_host\_name** (STRING): Host name of the local node.  
**localnodeid** (LONGINT): Node ID of the local node.  
**partner\_host\_name** (STRING): Host name of the partner node.  
**partnernodeid** (LONGINT): Node ID of the partner node.  
**reason** (STRING): Reason takeover was disabled.

## clam.update events

### clam.update.partner.state

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) updates the failover state of the partner.

#### Corrective Action

(None).

#### Syslog Message

CLAM on node (ID=%llu) updated failover state of partner (ID=%llu) to %s.

## Parameters

**localnodeid** (LONGINT): Node ID of the local node.  
**partnernodeid** (LONGINT): Node ID of the partner node.  
**partner\_state** (STRING): Description of the partner state.

# clam.valid events

## clam.valid.config

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services (KCS) Connectivity, Liveliness, Availability Monitor (CLAM) determines that the local node is in a valid configuration for providing CLAM functionality.

### Corrective Action

No action is needed.

### Syslog Message

Local node (name=%s, id=%llu) is operating in a suitable configuration for providing CLAM functionality.

## Parameters

**localhostname** (STRING): Host name of the local node.  
**localnodeid** (LONGINT): Node ID of the local node.  
**partnerhostname** (STRING): Host name of the partner node.  
**partnernodeid** (LONGINT): Node ID of the partner node.

# clam.voting events

## clam.voting.state.change

### Severity

INFORMATIONAL

### Description

This message occurs when the Kernel Cluster Services Connectivity, Liveliness, Availability Monitor (CLAM) detects that the voting state of the node changed.

### Corrective Action

(None).

### Syslog Message

CLAM on node (name=%s, ID=%llu) has %s voting for its HA partner node(name=%s, ID=%llu).

## Parameters

**local\_hostname** (STRING): Host name of the local node.  
**localId** (LONGINT): Node ID of the local node.  
**state** (STRING): New voting state of the node.  
**partner\_hostname** (STRING): Host name of the HA partner node.



**partnerId** (LONGINT): Node ID of the HA partner node.

# clone events

## clone.split events

### clone.split.objcopy.failed

#### Severity

ERROR

#### Description

This message occurs when the volume clone split object copy operation (which copies objects from source volume to a clone) fails because it is unable to read source objects due to unauthorized access to the object store container. The volume clone split object copy operation will be aborted and must be restarted manually after a valid access token is added to the object store configuration.

#### Corrective Action

Ensure that you have added the shared access signature token to the object store configuration. The access token can be added to the configuration using the "storage aggregate object-store config modify" command. Additionally, ensure that the configuration of your object store has not changed, and that login and connectivity information are still accurate. The volume clone split object copy operation can be restarted using the "volume clone split start -match-parent-storage-tier true" command. If the condition persists, contact NetApp technical support.

#### Syslog Message

Clone split object copy operation to copy objects from source volume to clone failed on volume "%s" (aggregate %s) because of unauthorized access to the object store container. Aborting the volume clone object copy split operation.

#### Parameters

**volume\_name** (STRING): Clone volume.

**aggregate\_name** (STRING): FabricPool aggregate.

# cloud events

## cloud.aws events

### cloud.aws.iamCredsExpired

#### Severity

ERROR

#### Description

This message occurs when the IAM Role thread acquires Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials that have already expired. The credentials are acquired from the AWS metadata server using the IAM role, and are used in signing API requests to AWS S3.

#### Corrective Action

Log in to the AWS EC2 Management Console, click on the 'Instances' page, and then find and check the health of the instance pertaining to this ONTAP Cloud deployment. Also verify that the AWS IAM role associated with the instance exists and grants proper privileges to the instance.

#### Syslog Message

ONTAP acquired AWS credentials associated with the AWS IAM role named '%s' on the AWS metadata server that have already expired. Node %s.

#### Parameters

**iamRole** (STRING): Name of the IAM role associated with this ONTAP Cloud® instance.  
**nodeUuid** (STRING): UUID of the node.

### cloud.aws.iamCredsInvalid

#### Severity

ERROR

#### Description

This message occurs when the system acquires Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials that are not valid.

#### Corrective Action

In the AWS EC2 Management Console, find and check the health of the instance pertaining to this ONTAP Cloud deployment. Also, verify that the AWS IAM role associated with the instance exists and grants proper privileges to the instance.

#### Syslog Message

AWS credentials that were acquired by ONTAP are not valid. AWS IAM role: %s. Node: %s.

#### Parameters

**iamRole** (STRING): Name of the IAM role associated with this ONTAP Cloud instance.  
**nodeUuid** (STRING): UUID of the node.

## cloud.aws.iamCredsNotFound

### Severity

ERROR

### Description

This message occurs when the cloud credentials thread cannot acquire Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the AWS metadata server. The credentials are used in signing API requests to AWS S3.

### Corrective Action

Log in to the AWS EC2 Management Console, click on the 'Instances' page, and then find and check the health of the instance pertaining to this Cloud Volumes ONTAP® deployment. Also verify that the AWS IAM role associated with the instance exists and grants proper privileges to the instance.

### Syslog Message

ONTAP cannot acquire credentials associated with the AWS IAM role or GCP Service Account named '%s' on the metadata server. Node %s.

### Parameters

**iamRole** (STRING): Name of the AWS IAM role associated associated with this Cloud Volumes ONTAP® instance.

**nodeUuid** (STRING): UUID of the node.

## cloud.aws.iamNotInitialized

### Severity

NOTICE

### Description

This message occurs when a module attempts to access Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the cloud credentials thread before the cloud thread is finished initializing.

### Corrective Action

Wait for the cloud credential thread, and by extension the system, to complete initialization.

### Syslog Message

A module attempted to access credential information before the cloud credential thread initialized on node %s.

### Parameters

**nodeUuid** (STRING): UUID of the node.

## cloud.aws.iamRoleInvalid

### Severity

ERROR

### Description

This message occurs when the system acquires an Amazon Web Services (AWS) Identity and Access Management (IAM) role on the AWS metadata server that is not valid.

### Corrective Action

In the AWS EC2 Management Console, find and check the health of the instance pertaining to this ONTAP Cloud deployment. Also, verify that the AWS IAM role associated with the instance exists and grants proper privileges to the instance.

### Syslog Message

IAM role "%s" on the AWS metadata server is not valid for this ONTAP instance (node %s).

### Parameters

**iamRole** (STRING): Name of the IAM role associated with this ONTAP Cloud instance.  
**nodeUuid** (STRING): UUID of the node.

## cloud.aws.iamRoleNotFound

### Severity

ERROR

### Description

This message occurs when the IAM Role thread cannot find an Amazon Web Services (AWS) Identity and Access Management (IAM) role on the AWS metadata server. The IAM role is needed to obtain role-based credentials used in signing API requests to AWS S3.

### Corrective Action

Log in to the AWS EC2 Management Console, click on the 'Instances' page, and then find and check the health of the instance pertaining to this ONTAP Cloud® deployment. Also verify that the AWS IAM role associated with the instance exists.

### Syslog Message

ONTAP node %s cannot find an IAM role on the AWS metadata server.

### Parameters

**nodeUuid** (STRING): UUID of the node.

## cloud.aws.metadataConnFail

### Severity

ERROR

### Description

This message occurs when the IAM Role thread cannot establish a communication link with the Amazon Web Services (AWS) metadata server. Communication must be established to acquire the necessary AWS Identity and Access Management (IAM) role-based credentials used in signing API requests to AWS S3.

### Corrective Action

Log in to the AWS EC2 Engine Management Console, click on the 'Instances' page, and then find and check the health of the instance pertaining to this Cloud Volumes ONTAP® deployment.

### Syslog Message

ONTAP cannot communicate with the AWS metadata server needed to acquire IAM role-based or Service Account based credentials for access to AWS S3 Storage. Error returned: %s Node: %s

## Parameters

**error** (STRING): Error code and string, if applicable.

**nodeUuid** (STRING): UUID of the node.

# cloud.azurecredsinvalid events

## cloud.azureCredsInvalid

### Severity

ERROR

### Description

This message occurs when the system acquires Azure Managed Service Identity (MSI) credentials that are not valid.

### Corrective Action

In the Azure Portal, find and check the health of the instance pertaining to this Cloud Volumes ONTAP® deployment. Also, verify that the MSI associated with the instance is enabled and grants proper privileges to the instance.

### Syslog Message

Azure MSI credentials that were acquired by ONTAP are not valid. Node UUID: %s.

## Parameters

**nodeUuid** (STRING): UUID of the node.

# cloud.bin events

## cloud.bin.objcopy.failed

### Severity

ERROR

### Description

This message occurs when the cloud bin object copy operation (which copies objects from a source object store to a destination object store) fails because it is unable to read source objects due to unauthorized access to the object store container.

### Corrective Action

Ensure that you have added the shared access signature token to the source object store configuration. The access token can be added to the configuration using the "storage aggregate object-store config modify" command. Additionally, ensure that the configuration of your object store has not changed, and that login and connectivity information are still accurate. If the condition persists, contact NetApp technical support.

### Syslog Message

Cloud bin object copy operation to copy objects from a source object store to destination object store failed for "%s" operation because of unauthorized access to the source object store container ("%s").

## Parameters

**op\_name** (STRING): Operation context in which object copy is called. This can be either clone split or svm migrate.

**container\_name** (STRING): Source object store container name

# cloud.creds events

## cloud.creds.metadataConnFail

### Severity

ERROR

### Description

This message occurs when the system cannot establish a communication link with the cloud provider's metadata server, in order to acquire the necessary temporary credentials used in signing API requests to the cloud object store server.

### Corrective Action

Log in to the cloud provider's management console to examine the current set of running instances. Find the instance pertaining to this deployment of Cloud Volumes ONTAP® deployment, and then check its health from the console.

### Syslog Message

ONTAP cannot communicate with the cloud provider's metadata server needed to acquire the temporary credentials used to access an object store server. Command: %s Error returned: %s Node: %s

## Parameters

**command** (STRING): Command used to contact the metadata server.

**error** (STRING): Error text received as response to the metadata server command.

**nodeUuid** (STRING): UUID of the node.

## cloud.creds.notFound

### Severity

ERROR

### Description

This message occurs when the system cannot acquire the temporary credentials from the cloud provider's metadata server that are necessary for signing API requests to the object storage server.

### Corrective Action

Log in to the management console of the cloud provider to examine the current set of running instances. Find and check the health of the instance pertaining to this Cloud Volumes ONTAP deployment. Also verify that the role associated with the instance exists and grants the proper privileges.

### Syslog Message

ONTAP cannot acquire credentials associated with the instance named '%s' on the metadata server. Node %s.

## Parameters

**role** (STRING): Name of the AWS IAM role or GCP Service Account or Azure Managed Service Identity associated with this Cloud Volumes ONTAP® instance.  
**nodeUuid** (STRING): UUID of the node.

## cloud.creds.notInitialized

### Severity

NOTICE

### Description

This message occurs when the system attempts to access credentials before initialization is complete.

### Corrective Action

(None).

### Syslog Message

Module on node %s attempted to access credential information before initialization was complete

### Parameters

**nodeUuid** (STRING): UUID of the node.

## cloud.gcpcredsinvalid events

### cloud.gcpCredsInvalid

### Severity

ERROR

### Description

This message occurs when the system acquires Google Cloud Platform (GCP) Service Account based credentials that are not valid.

### Corrective Action

In the Google Cloud Console, find and check the health of the instance pertaining to this Cloud Volumes ONTAP® deployment. Also, verify that the Service Account associated with the instance exists and grants proper privileges to the instance.

### Syslog Message

GCP credentials that were acquired by ONTAP are not valid. Node: %s.

### Parameters

**nodeUuid** (STRING): UUID of the node.



# clus events

## clus.net events

### clus.net.noNeighbors

#### Severity

ALERT

#### Description

This message occurs when the cluster network ports with "link" reported as "up" are neither receiving CDP nor LLDP neighbor announcements. This is not a supported configuration and can negatively impact the performance of the cluster.

#### Corrective Action

Use the "network port show -ipospace Cluster -link up" command to list the ports expected to receive CDP and/or LLDP announcements. Verify that CDP and/or LLDP is enabled on the switch the cluster ports are wired to. Use the "network device-discovery show" command to view the announcements for each of the cluster ports.

#### Syslog Message

The cluster ports with "link" reported as "up" are neither receiving CDP nor LLDP neighbor announcements.

#### Parameters

(None).

### clus.net.switchlessInv

#### Severity

NOTICE

#### Description

This message occurs when the "switchless-cluster" option is disabled, but one or more cluster ports are directly wired to a cluster port on another node. This could result in packet loss across the cluster network, causing file operations to fail as well as loss of quorum.

#### Corrective Action

Use the "network options switchless-cluster modify" command to enable support for switchless clusters.

#### Syslog Message

One or more cluster ports are directly wired to a cluster port on another node, but the "switchless-cluster" option is not enabled.

#### Parameters

(None).

### clus.net.switchlessMismatch

**Severity**

NOTICE

**Description**

This message occurs when the "switchless-cluster" option is enabled, but all of the cluster ports are wired to a switch. This configuration will negatively impact the performance of the cluster network.

**Corrective Action**

Use the "network options switchless-cluster modify" command to disable support for switchless clusters.

**Syslog Message**

All of the cluster ports are wired to a switch, but the "switchless-cluster" option is enabled.

**Parameters**

(None).

## clus.reboot events

### clus.reboot.internalError

**Severity**

ALERT

**Description**

This message occurs when the cluster reboot detection feature detects an internal software error due to insufficient system resources.

**Corrective Action**

Reboot the node by using the "system node reboot -node <nodename>" command. Contact NetApp technical support if further assistance is needed.

**Syslog Message**

Cluster reboot detection feature detected an internal software error. Error: '%s'.

**Parameters**

**error** (STRING): Internal software error.

### clus.reboot.noBootSequence

**Severity**

ERROR

**Description**

This message occurs when the cluster reboot detection feature fails to obtain the node boot sequence from the /var/BOOT\_SEQ file.

**Corrective Action**

Contact NetApp technical support.

## Syslog Message

Cluster reboot detection feature failed to obtain the node boot sequence from /var/BOOT\_SEQ file.

## Parameters

(None).

## clus.reboot.persistencyError

### Severity

ERROR

### Description

This message occurs when the cluster reboot detection feature fails to save the node reboot information persistently.

### Corrective Action

Reboot the node by using the "system node reboot -node nodename" command. Contact NetApp technical support if further assistance is needed.

## Syslog Message

Cluster reboot detection feature failed to save the node reboot information persistently. Error: '%s'.

## Parameters

**error** (STRING): Error returned by the save operation.

## clus.reboot.rebootDetected

### Severity

NOTICE

### Description

This message occurs when a cluster reboot is detected.

### Corrective Action

(None).

## Syslog Message

Cluster reboot was detected at %s. Sequence number for the cluster reboot was %d. Estimated beginning time of the cluster reboot was %s.

## Parameters

**endtime** (STRING): End time of the cluster reboot.

**sequence** (LONGINT): Sequence number of the detected cluster reboot.

**begintime** (STRING): Estimated beginning time of the cluster reboot based on the last known time the cluster was operational.

# cluster events

## cluster.epsilon events

### cluster.epsilon.assign.fail

#### Severity

ERROR

#### Description

This message occurs when automatic assignment of epsilon to a node fails, alerting the customer that there is no epsilon on their system. Attempts to assign epsilon continue until automatic assignment succeeds or epsilon is assigned manually.

#### Corrective Action

Try to manually assign epsilon to a healthy node using the command 'cluster modify -node node-name -epsilon true'. Contact NetApp technical support if epsilon cannot be assigned.

#### Syslog Message

Epsilon could not be automatically assigned to node %s with error %s.

#### Parameters

**node** (STRING): Name of the node that epsilon was being assigned to.  
**error** (STRING): Error string detailing why epsilon assignment failed.

### cluster.epsilon.assigned

#### Severity

INFORMATIONAL

#### Description

This message occurs when epsilon is assigned to a node either manually by a user-initiated command or automatically by the system.

#### Corrective Action

(None).

#### Syslog Message

Epsilon was assigned to node %s by %s.

#### Parameters

**node** (STRING): Name of the node epsilon was assigned to.  
**actor** (STRING): Whether the change was made manually by a user command or automatically by the system.

### cluster.epsilon.removed

#### Severity

INFORMATIONAL

## Description

This message occurs when epsilon is removed from a node either manually by a user-initiated command or automatically by the system.

## Corrective Action

(None).

## Syslog Message

Epsilon was removed from node %s by %s.

## Parameters

**node** (STRING): Name of the node that epsilon was removed from.

**actor** (STRING): Whether the change was made manually by a user command or automatically by the system.

# cluster.ha events

## cluster.ha.not.configured

### Severity

ERROR

### Description

This message occurs when cluster HA has not been configured on a 2-node cluster with HA mode enabled.

### Corrective Action

Use the 'cluster ha modify -configured true' command to configure cluster HA.

### Syslog Message

Cluster HA has not been configured. Cluster HA must be configured on a 2-node cluster to ensure data access availability in the event of storage failover.

### Parameters

(None).

# cluster.image events

## cluster.image.update.aborted

### Severity

INFORMATIONAL

### Description

This message occurs when a non-disruptive data ONTAP upgrade of the cluster was aborted by the user.

### Corrective Action

(None).

### **Syslog Message**

A non-disruptive data ONTAP upgrade of the cluster was aborted by the user

### **Parameters**

(None).

## **cluster.network events**

### **cluster.network.config**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a two node cluster transitions between a switched and switchless configuration.

#### **Corrective Action**

(None).

### **Syslog Message**

The cluster network configuration transitioned to %s.

### **Parameters**

**current\_configuration** (STRING): Displays if the current configuration is switched or switchless.

## **cluster.op events**

### **cluster.op.fail**

#### **Severity**

ERROR

#### **Description**

This message occurs when a "cluster create" or a "cluster join" or a "cluster unjoin" operation is attempted but does not succeed.

#### **Corrective Action**

Run the command specified to complete the cluster operation. Resolve the issue, and then try the command again.

### **Syslog Message**

(None).

### **Parameters**

**operation\_type** (STRING): Type of cluster operation.

**node** (STRING): Name of the node that attempted the cluster operation.

**serial** (STRING): System serial number of the node that attempted the cluster operation.

## cluster.op.not.complete

### Severity

ERROR

### Description

This message occurs when a cluster create or cluster join operation fails to complete. The node is not fully operational and many commands will not work properly until this is resolved.

### Corrective Action

Run the command specified to complete the cluster operation. Resolve the issue, and then try the command again.

### Syslog Message

Cluster operation (%s) failed to complete for node %s. The node is not fully operational.

### Parameters

**command** (STRING): Command for the cluster operation.

**node** (STRING): Name of the node.

## cluster.op.rerun.fail

### Severity

ERROR

### Description

This message occurs when a previously failed cluster create or cluster join operation that is automatically rerun in background fails. The operation is automatically retried a maximum of 'max\_retry\_count' times at an interval of 'retry\_interval' seconds.

### Corrective Action

To view the reason for failure, run the operation manually. Alternatively, allow the automatic reruns to execute in case the issue is temporary or fixed. To view the progress of these background attempts, use the command 'event log show -event cluster.op.rerun\*'. You may rerun the operation manually at any time, without waiting for these background runs to complete.

### Syslog Message

Attempt to %s cluster named '%s' in the background failed. This was retry attempt %d of %d.

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**cluster\_name** (STRING): Name of the cluster that is being created or joined.

**retry\_count** (INT): Number of attempts made to rerun the cluster create or join operations in the background.

**max\_retry\_count** (INT): Maximum number of retry attempts allowed to rerun the cluster create or join operation in the background.

**retry\_interval** (INT): Interval of time, in seconds, between each rerun attempt.

## cluster.op.rerun.start

## Severity

INFORMATIONAL

## Description

This message occurs when a previously failed cluster create or cluster join operation is rerun in the background in an attempt to successfully complete the operation. The operation is automatically retried a maximum of 'max\_retry\_count' times at an interval of 'retry\_interval' seconds.

## Corrective Action

To view the reason for failure, run the operation manually. Alternatively, allow the automatic reruns to execute in case the issue is temporary or fixed. To view the progress of these background attempts, use the command 'event log show -event cluster.op.rerun\*' You may rerun the operation manually at any time, without waiting for these background runs to complete.

## Syslog Message

Attempt to %s cluster named '%s' in the background. This is retry attempt %d of %d.

## Parameters

**operation\_type** (STRING): Type of cluster operation.

**cluster\_name** (STRING): Name of the cluster that is being created or joined.

**retry\_count** (INT): Number of attempts to retry the cluster create or join operation in the background.

**max\_retry\_count** (INT): Maximum number of rerun attempts allowed.

**retry\_interval** (INT): Interval of time, in seconds, between each rerun attempt.

## cluster.op.rerun.success

### Severity

INFORMATIONAL

### Description

This message occurs when a previously failed cluster create or cluster join operation automatically rerun in the background is successful.

### Corrective Action

(None).

### Syslog Message

Cluster %s operation run in the background completed successfully for cluster name '%s'.

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**cluster\_name** (STRING): Name of the cluster that is being created or joined.

## cluster.op.start

### Severity

INFORMATIONAL

### Description

This message occurs when a "cluster create" or a "cluster join" or a "cluster unjoin" operation is run.



### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**node** (STRING): Name of the node being operated on or empty if the name is not yet set.

**serial** (STRING): System serial number of the node being operated on or empty if the serial number is not set.

## cluster.op.success

### Severity

INFORMATIONAL

### Description

This message occurs when a cluster has been successfully created or a node has successfully joined or unjoined the cluster.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**node** (STRING): Name of the node that successfully performed the cluster operation.

**serial** (STRING): System serial number of the node that successfully performed the cluster operation.

## cluster.port events

### cluster.port.unsupport.speed

### Severity

ERROR

### Description

This message occurs when the cluster ports are operating at an unsupported speed.

### Corrective Action

Ensure that the cluster port has a supported network interface card installed, and that it is operating at the supported speed.

### Syslog Message

%s node %s cluster ports are not operating at a supported speed. Ensure that ports %s have supported network interface card installed, and are operating at a supported speed: %s.

## Parameters

- model\_name** (STRING): Model on which the misconfigured ports reside.
- node** (STRING): Node on which the misconfigured ports reside.
- portlist** (STRING): List of misconfigured ports.
- supported\_speed** (STRING): Speed at which the port is supported.

# cluster.repl events

## cluster.repl.localNW.unreachable

### Severity

NOTICE

### Description

This message occurs when the cluster replication heartbeat to a peer cluster cannot reach the local management gateway process for a minimum of 300 seconds. This potentially affects MetroCluster(tm) and Vserver disaster recovery (DR) applications.

### Corrective Action

(None).

### Syslog Message

The cluster replication heartbeat could not reach the local management gateway process for at least %d seconds.

### Parameters

- durationSinceLastSuccess** (INT): Number of seconds since the problem first occurred.

## cluster.repl.peer.reachable

### Severity

NOTICE

### Description

This message occurs when the cluster replication heartbeat can now reach the peer cluster, after previous failures resulted in raising one of the following events: cluster.repl.localNW.unreachable, cluster.repl.peerNW.unreachable, or cluster.repl.remoteApp.unreachable.

### Corrective Action

(None).

### Syslog Message

The cluster replication heartbeat can now reach peer cluster %s.

### Parameters

- peerIPAddress** (STRING): Peer cluster IP address that can now be reached.

## cluster.repl.peerNW.unreachable

## Severity

ALERT

## Description

This message occurs when the cluster replication heartbeat to the peer cluster fails for a minimum of 300 seconds. This potentially affects MetroCluster(tm) and Vserver disaster recovery (DR) applications.

## Corrective Action

Verify network connectivity between the local cluster and the peer cluster. Check the cluster peer EMS messages for more information.

## Syslog Message

The cluster replication heartbeat to peer cluster %s failed for at least %d seconds.

## Parameters

**peerIPAddress** (STRING): Cluster peer IP address that could not be reached.

**durationSinceLastSuccess** (INT): Number of seconds since the problem first occurred.

## cluster.repl.remoteApp.unreachable

### Severity

ALERT

### Description

This message occurs when the cluster replication heartbeat to the remote cluster replication application fails for a minimum of 300 seconds. This potentially affects MetroCluster(tm) and Vserver disaster recovery (DR) applications.

### Corrective Action

Check the state of the replication application process running on the peer cluster by issuing the 'cluster ring show -unitname crs' command.

### Syslog Message

The cluster replication heartbeat to the remote application running on peer cluster %s failed for at least %d seconds.

### Parameters

**peerIPAddress** (STRING): Remote application cluster IP address that could not be reached.

**durationSinceLastSuccess** (INT): Number of seconds since the problem first occurred.

## cluster.switch events

### cluster.switch.count

### Severity

INFORMATIONAL

### Description

Redundant configuration is missing for cluster switches. If the remaining cluster switch fails, you can no longer access the cluster.

## Corrective Action

1) Ensure that a redundant pair of cluster switches exists in the network. 2) Ensure that cluster ports are connected to both switches. 3) Enable the Cisco Discovery Protocol (CDP)/Industry Standard Discovery Protocol (ISDP) on the switches, if you previously disabled it. CDP/ISDP is enabled by default. Refer to your switch documentation for instructions. 4) If Data ONTAP cannot automatically discover a switch, use the "system cluster-switch create" command to configure discovery and monitoring of the switch.

## Syslog Message

Cluster switch count: %d but requires %d.

## Parameters

**count** (INT): Number of cluster switches discovered.  
**required** (INT): Number of cluster switches required.

## cluster.switch.pwr.count

### Severity

INFORMATIONAL

### Description

The power supply on the switch is missing. The remaining power supply is a single point of failure. The switch might shut down with the next PSU failure.

## Corrective Action

1) Check whether the switch's power supply unit is properly inserted into the chassis of the switch.

## Syslog Message

Cluster switch: %s has: %d power supplies and requires %d.

## Parameters

**switch\_name** (STRING): Cluster switch name  
**count** (INT): The switch power supply count found.  
**required** (INT): The switch power supply count required.

## cluster.switch.pwr.fail

### Severity

INFORMATIONAL

### Description

The power supply on the switch is missing or is not operational.

## Corrective Action

1) Ensure that the power supply mains supplying power to the switch is turned on. 2) Ensure that the power cord is connected to the power supply of the switch. 3) Contact technical support if the alert persists.

## Syslog Message

Cluster switch: %s power supply: %s status: %s.

## Parameters

**switch\_name** (STRING): Cluster switch name

**pwr\_supply\_name** (STRING): The power supply name that is failing.  
**status** (STRING): The power supply status.

## cluster.switch.pwr.non.operational

### Severity

INFORMATIONAL

### Description

The power supply on the switch is not operational.

### Corrective Action

1) Ensure that the power supply unit is turned on through switch CLI. 2) Run diagnostics on the switch to check the operational status of the power supply unit (PSU). For instructions about how to run switch diagnostics, see the switch's configuration guide. 3) Contact technical support if the alert persists.

### Syslog Message

Cluster switch: %s power supply: %s status: %s.

### Parameters

**switch\_name** (STRING): Cluster switch name  
**pwr\_supply\_name** (STRING): The power supply name that is failing.  
**status** (STRING): The power supply status.

## cluster.switch.reboot

### Severity

INFORMATIONAL

### Description

The cluster switch has recently rebooted. Communication problems and cluster connectivity issues might occur.

### Corrective Action

1) If the cluster switch was not rebooted on purpose, then check the switch to ensure that it is operating normally. 2) If errors persist, contact technical support for further assistance.

### Syslog Message

Cluster switch: %s rebooted at: %ld

### Parameters

**switch\_name** (STRING): Cluster switch name  
**time** (INT): Last reported time before reboot.

## cluster.switch.unsupported

### Severity

INFORMATIONAL

### Description

Unsupported cluster switch is detected. Communication problems and cluster connectivity issues occur.

### Corrective Action

Unsupported cluster switch found. Connect a supported cluster switch to cluster network.

### Syslog Message

Cluster switch: %s model: %s not supported.

### Parameters

**switch\_name** (STRING): Cluster switch name

**model** (STRING): The switch model.

## cluster.vol events

### cluster.vol.create.success

#### Severity

NOTICE

#### Description

This message occurs when the volume required to store the system data needed for cluster-wide storage is created successfully after one or more attempts to create it failed. An event was previously raised for the failures.

#### Corrective Action

(None).

#### Syslog Message

Successfully created the volume for feature %s, after earlier attempts to create it failed.

#### Parameters

**Feature** (STRING): Feature using the cluster-wide storage.

### cluster.vol.delete.failed

#### Severity

ALERT

#### Description

This message occurs when the volume that was used for cluster-wide storage cannot be deleted.

#### Corrective Action

Study the error message and take steps to rectify the problem. For example, if the error message indicates that the volume delete operation failed because the aggregate is offline, use the "aggr online -aggregate aggregate\_name" command to bring it back online. If the error message does not indicate an obvious problem, contact NetApp technical support.

#### Syslog Message

Could not delete volume %s %d times on aggregate %s for feature %s. Most recent delete before this EMS failed with error message: "%s".

## Parameters

**Volume** (STRING): Volume name.

**AttemptsMade** (INT): Count of failed attempts to delete the volume.

**Aggregate** (STRING): Aggregate name.

**Feature** (STRING): Feature using the cluster-wide storage.

**ErrorMessage** (STRING): Error message resulting from the delete failure at the time of EMS generation.

## cluster.vol.delete.success

### Severity

NOTICE

### Description

This message occurs when the volume that was used for cluster-wide storage is deleted successfully after one or more attempts to delete it failed. An event was previously raised for the failures.

### Corrective Action

(None).

### Syslog Message

Successfully deleted volume %s on aggregate %s for feature %s, after earlier attempts to delete it failed.

## Parameters

**Volume** (STRING): Volume name.

**Aggregate** (STRING): Aggregate name.

**Feature** (STRING): Feature using the cluster-wide storage.

## cluster.vol.feature.inoper

### Severity

ALERT

### Description

This message occurs when it is detected that a feature relying on cluster-wide storage has been unable to create its storage. This condition leaves the feature inoperative.

### Corrective Action

Look for any instances of the callhome.clus.vol.cre.fail event and follow any corrective action suggested there. For more information or assistance, contact NetApp technical support.

### Syslog Message

Feature %s is inoperative due to failures creating the cluster-wide storage that it requires.

## Parameters

**Feature** (STRING): Feature using the cluster-wide storage.

## cluster.vol.feature.ok

### Severity

INFORMATIONAL

## Description

This message occurs when it is detected that a feature relying on cluster-wide storage has been able to create its storage after previously not being able to.

## Corrective Action

(None).

## Syslog Message

Feature %s degraded state has been cleared.

## Parameters

**Feature** (STRING): Feature using the cluster-wide storage.

## cluster.vol.on.same.aggr

## Severity

ALERT

## Description

This message occurs when both active and standby volumes used for cluster-wide storage are hosted on the same aggregate. This is typically due to there not being a suitable second aggregate to host the standby volume. In this configuration, the loss of the single aggregate would remove all cluster-wide storage for the feature, which would severely affect its functionality.

## Corrective Action

Use the "volume show" command to identify the size of the specified volumes and on which aggregate they are hosted. Use the "storage aggregate show -fields mirror" command to identify if the aggregate that currently hosts the volumes is mirrored. Use the "storage aggregate show -fields mirror,size,usedsize,state" command to find an online aggregate with the same mirror attribute and with enough free space to host the standby volume. If a suitable aggregate cannot be found, then bring an offline aggregate back online using the "storage aggregate online" command or create a new aggregate using the "storage aggregate create" command. If creating a new aggregate, ensure that it is large enough to host the volume and that it has the same mirror attribute as the current aggregate. After a suitable aggregate has been identified or created, move the standby volume to this aggregate using the "volume move start" command. For more information or assistance, contact NetApp technical support.

## Syslog Message

Both active and standby cluster-wide storage volumes (%s, %s) for %s feature are on the same aggregate.

## Parameters

**ActiveVolume** (STRING): Active volume name.

**StandbyVolume** (STRING): Standby volume name.

**Feature** (STRING): Feature using the cluster-wide storage.



# cmd events

## cmd.result events

### cmd.result.reboot

**Severity**

NOTICE

**Description**

This message occurs when the system reboots as a consequence of executing the specified command.

**Corrective Action**

(None).

**Syslog Message**

System rebooted for ('%s').

**Parameters**

**cmd** (STRING): Command for which the system rebooted.

## cmd.sasadmin events

### cmd.sasadmin.fw\_update.fail

**Severity**

ERROR

**Description**

This message occurs when an adapter firmware update failure occurs.

**Corrective Action**

Verify that the firmware file is valid and try updating the firmware again. If the firmware update still fails, contact NetApp technical support.

**Syslog Message**

Adapter firmware update failed for adapter: %s.

**Parameters**

**adapterName** (STRING): Name of the associated SAS host bus adapter.

## cmd.shelfchk events

### cmd.shelfchk.config.OK

**Severity**

INFORMATIONAL

**Description**

This event is generated while verifying disk shelf cabling and configuration. The operator indicates that the correct disk LEDs have illuminated.

**Corrective Action**

(None).

**Syslog Message**

Disk Configuration - No Errors Identified.

**Parameters**

(None).

**cmd.shelfchk.config.warning****Severity**

ERROR

**Description**

This event is generated while verifying disk shelf cabling and configuration on a channel. The operator has indicated that the intended set of disk LEDs is not illuminated.

**Corrective Action**

Verify that the adapters and storage shelves are cabled correctly.

**Syslog Message**

Disk Configuration - Failure Identified by Operator.

**Parameters**

(None).

**cmd.start events****cmd.start.unsuccessful****Severity**

ERROR

**Description**

This message occurs when the Counter Manager daemon (cmd) fails to start. The failure impacts the continuous archiving of performance data, which is used by technical support to troubleshoot system issues.

**Corrective Action**

The startup failure might be due to improper permissions for the '/mroot/etc/log' directory, or a full disk. Verify that the permissions of the '/mroot/etc/log' directory are 'rwxr-xr-x' by executing the 'fsecurity show /vol/vol0/etc/log' command from the nodeshell. Verify that the '/mroot' file system is not full. After fixing the cause of the failure, restart the cmd by running the 'spmctl -e -h cmd' command from the freebsd shell.

**Syslog Message**

The cmd startup failed.

**Parameters**

(None).

# cmds events

## cmds.config events

### cmds.config.file.delete

**Severity**

NOTICE

**Description**

This message occurs when there is an error in deleting the specified file.

**Corrective Action**

(None).

**Syslog Message**

%s: Error in deleting %s: (%s).

**Parameters**

**funcname** (STRING): Function name.

**fname** (STRING): File name.

**errMessage** (STRING): Error message.

## cmds.software events

### cmds.software.dirErr

**Severity**

ERROR

**Description**

This message occurs because the system cannot create the directory at the designated path.

**Corrective Action**

(None).

**Syslog Message**

Could not create directory %s: %s

**Parameters**

**path** (STRING): Path where directory was supposed to be created.

**errorMsg** (STRING): Error message.

### cmds.software.installDone

**Severity**

INFORMATIONAL

**Description**

This message occurs during software update and installation to inform you that the installation is complete.

**Corrective Action**

(None).

**Syslog Message**

Software: Installation of %s was completed.

**Parameters**

**fileName** (STRING): Name of the file to be updated/installed.

**cmds.software.installNotDone****Severity**

INFORMATIONAL

**Description**

This message occurs during software update and installation to inform you that installation failed.

**Corrective Action**

(None).

**Syslog Message**

Software: Installation of %s failed.

**Parameters**

**fileName** (STRING): Name of the file to be updated/installed.

**cmds.sysconf events****cmds.sysconf.haltERR1****Severity**

ALERT

**Description**

This message occurs to log the reason for system halt, before actually system is halted.

**Corrective Action**

(None).

**Syslog Message**

Halting system due to %d configuration problems.

**Parameters**

**errCount** (INT): Number of error which caused the system halt.

## **cmds.sysconf.haltERR2**

### **Severity**

ALERT

### **Description**

This message occurs to log the reason for system halt, before actually system is halted.

### **Corrective Action**

(None).

### **Syslog Message**

Halting system due to a configuration problem.

### **Parameters**

(None).

## **cmds.sysconf.inconstRule**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs because the rule for the card is applicable to both non\_cluster\_only and cluster\_only systems, which is inconsistent.

### **Corrective Action**

(None).

### **Syslog Message**

Rule for model %s card %s is inconsistent: it applies both to cluster-only and non-cluster-only.

### **Parameters**

**model** (STRING): Model ID.

**card** (STRING): Card ID.

## **cmds.sysconf.logErr**

### **Severity**

NOTICE

### **Description**

This message occurs to log the error through syslog as an Error or Informational message depending upon the error count which is in this case greater than 0.

### **Corrective Action**

(None).

### **Syslog Message**

%s.

## Parameters

**errMsg** (STRING): Error message to be logged.

## cmds.sysconf.logInfo

### Severity

INFORMATIONAL

### Description

This message occurs to log the error through syslog as an Error or informational message depending upon the error count which is in this case less than or equal to 0.

### Corrective Action

(None).

### Syslog Message

%s.

## Parameters

**errMsg** (STRING): Error message to be logged.

## cmds.sysconf.perIntSmall

### Severity

INFORMATIONAL

### Description

This message occurs because the value for a periodic interval for internal testing is set too small, less than the minimum allowed.

### Corrective Action

(None).

### Syslog Message

options.sysconfig.periodic\_interval of %d minutes is too small; using %d.

## Parameters

**valueUsed** (INT): Value used for testing purpose.

**defaultMinValue** (INT): Minimum value allowed for a periodic interval.

## cmds.sysconf.syslogger

### Severity

NOTICE

### Description

This message occurs to send the output to syslog rather than command prompt because the option '-l' with sysconfig command is not documented and is included in DEBUG builds only.

**Corrective Action**

(None).

**Syslog Message**

%s.

**Parameters**

**message** (STRING): 'sysconfig -l' Command output message.



# cmpr events

## cmpr.algo events

### cmpr.algo.not.enabled

#### Severity

ALERT

#### Description

This message occurs when compression algorithm specified has not been enabled on the node.

#### Corrective Action

Contact technical support.

#### Syslog Message

Volume '%s' has '%s' compression enabled but the necessary bootarg(s) have not been set on the node.

#### Parameters

**vol** (STRING): Volume name.

**comp\_algo** (STRING): Compression Algorithm.

# cna events

## cna.dump events

### cna.dump

#### Severity

NOTICE

#### Description

This message occurs when an unexpected event or illegal condition is detected by the Converged Network Adapter (CNA) driver and the contents of the adapter's static RAM and memory must be dumped. After the dump, the adapter is reset and the contents of the dump are stored in a file in the `/etc/log/` directory.

#### Corrective Action

No corrective action is needed; the adapter was reset.

#### Syslog Message

CNA coredump generated for CNA adapter CNA\_%, %s

#### Parameters

**adapter** (STRING): Name of the Converged Network Adapter.

**reason** (STRING): Reason for the dump.

### cna.dump.saved

#### Severity

NOTICE

#### Description

This message occurs when an unexpected event or illegal condition is detected by the Converged Network Adapter (CNA) driver and the contents of the adapter's static RAM and memory are saved. The dump files are stored on the appliance's root volume in the `/etc/log/` directory, with the following file name format:  
`[date][time].bin`

#### Corrective Action

Send the dump file to technical support for analysis.

#### Syslog Message

A dump for CNA adapter CNA\_%, %s was stored in %s

#### Parameters

**adapter** (STRING): Name of the CNA.

**filename** (STRING): Name of the file containing the dump.

## cna.initfailed events

## **cna.initFailed**

### **Severity**

ALERT

### **Description**

This message occurs when the converged network adapter (CNA) fails to initialize.

### **Corrective Action**

Take corrective action based on the reason indicated for the failure.

### **Syslog Message**

Initialization of adapter %s failed, reason: %s.

### **Parameters**

**adapter** (STRING): Name of the Adapter.

**reason** (STRING): Reason for the failure.

## **cna.sfp events**

### **cna.sfp.notSupported**

### **Severity**

ERROR

### **Description**

This message occurs when the converged network adapter (CNA) detects an unsupported small form-factor Pluggable (SFP) in the adapter.

### **Corrective Action**

Replace the unsupported SFP with a supported SFP.

### **Syslog Message**

The SFP in adapter %s is not supported.

### **Parameters**

**adapter** (STRING): Name of the adapter.

# command events

## command.complex events

### command.complex.query

#### Severity

ERROR

#### Description

This message occurs when a custom role associated with a Storage Virtual Machine has a command directory with complex query.

#### Corrective Action

Remove the complex query. See [BUG 915877](#) for more information.

#### Syslog Message

The %s command directory in %s role under Storage Virtual Machine %s, will not be available in merged role.

#### Parameters

**vserver** (STRING): Name of the Storage Virtual Machine.

**rolename** (STRING): Name of the role associated with the Virtual Storage Machine, with a command directory with a complex query.

**cmddirname** (STRING): Name of the command directory with complex query.

# config events

## config.atanotsupported events

### config.ATANotSupported

**Severity**

ERROR

**Description**

This message occurs when an ATA disk is found, but the appliance does not support it.

**Corrective Action**

Remove the ATA disks or upgrade your appliance to a model that supports ATA disks. Check your support provider's web site or knowledgebase for more information about appliance models that support ATA disks.

**Syslog Message**

ATA disks are not supported on this appliance. Disk %s is an ATA disk and should be removed.

**Parameters**

**disk** (STRING): Name of the disk.

## config.badlocalswitch events

### config.BadLocalSwitch

**Severity**

ALERT

**Description**

This message occurs when a disk is found to be attached to a switch, but there are also disks attached via local connections. Mixed switch-attached and direct-attached disks are an illegal configuration.

**Corrective Action**

Correct the cabling so that all disks are attached either directly or through switches. Consult the hardware cabling guides at the NetApp technical support site.

**Syslog Message**

Disk %s is attached to a switch. However, there are also local attached disks. This is an illegal configuration.

**Parameters**

**disk** (STRING): Name of the disk.

## config.badpoolassign events

### config.BadPoolAssign

## Severity

ERROR

## Description

This message occurs when SyncMirror® is in use and disks are assigned to different pools on the same interface or loop. This affects the fault isolation of the plexes in mirrored aggregates such that if a single interface fails, both plexes in an aggregate might fail or become degraded. If both plexes of an aggregate fail, the system enters multidisk panic.

## Corrective Action

If the disk is a spare disk, use the "storage disk assign -pool [0|1]" command to change the pool assignment of the disks on the interface so that they are assigned to only one pool. If the disk is a file system disk, use the "storage disk replace" command to replace the disk from the incorrect pool with a spare disk from the correct pool.

## Syslog Message

Disk %s is in %s and other disks on this loop/domain are in %s. Disks/interfaces need to be in separate pools for SyncMirror.

## Parameters

**disk** (STRING): Name of the disk that fails.  
**disk\_pool** (STRING): Name of the pool that the disk is in.  
**other\_pool** (STRING): Name of the other pool.

# config.badslot events

## config.BadSlot

### Severity

ALERT

### Description

This message occurs when a disk is found to be attached to an invalid Host Adapter (HA) controller slot.

### Corrective Action

Consult the NetApp technical support guides for the supported slots for adapters in the hardware platform. If the adapter is in a supported slot, contact NetApp technical support for further assistance.

### Syslog Message

Disk %s is attached to an invalid HA slot. This is an illegal configuration.

### Parameters

**disk** (STRING): Name of the disk.

# config.badswitchlocal events

## config.BadSwitchLocal

### Severity

ALERT

## Description

This message occurs when a disk is found to be attached locally, but there are also disks attached via switches. Mixed switch-attached and direct-attached disks are an illegal configuration.

## Corrective Action

Correct the cabling so that all disks are attached either directly or through switches. Consult the hardware cabling guides at the NetApp technical support site.

## Syslog Message

Disk %s is attached locally. However, there are also switch attached disks. This is an illegal configuration.

## Parameters

**disk** (STRING): Name of the disk.

# config.bridge events

## config.bridge.adapter.invalid

### Severity

ERROR

### Description

This message occurs when an FC-to-SAS bridge is attached to an unsupported FC adapter port. Only QLogic® 4-Gb/s or 8-Gb/s adapters are supported.

### Corrective Action

Attach the FC-to-SAS bridge to a supported FC adapter port on the controller.

### Syslog Message

FC-to-SAS bridge %s S/N [%s] is attached to unsupported FC adapter port %s.

### Parameters

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**adapterName** (STRING): Name of the associated FC host bus adapter.

## config.bridge.extra.port

### Severity

ERROR

### Description

This message occurs when both FC ports of an FC-to-SAS bridge are attached to the same controller.

### Corrective Action

Follow the defined Fabric or Stretch MetroCluster wiring guide.

### Syslog Message

Both FC ports of FC-to-SAS bridge %s S/N [%s] are attached to this controller.

## Parameters

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

## config.bridge.fw.mismatch

### Severity

ERROR

### Description

This message occurs when two or more attached FC-to-SAS bridges are running different firmware revisions.

### Corrective Action

For more information about downloading and installing updated firmware, check your support provider's web site or knowledgebase.

### Syslog Message

Two or more %s FC-to-SAS bridges are running different firmware revisions.

## Parameters

**bridgeModel** (STRING): Name of the FC-to-SAS bridge model.

## config.bridge.SAS.miswired

### Severity

ERROR

### Description

This message occurs when the SAS disk shelves behind an FC-to-SAS bridge are not wired in the correct order.

### Corrective Action

Follow the defined Fabric or Stretch MetroCluster wiring guide.

### Syslog Message

The SAS disk shelves behind FC-to-SAS bridge %s S/N [%s] are not correctly wired.

## Parameters

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

## config.bridge.SAS2Domain

### Severity

ERROR

### Description

This message occurs when a disk is multipathed to one SAS domain; that is, both ports of the disk are visible from the same bridge FC port.



### Corrective Action

Recable the disk shelf containing the specified disk, so that the two disk ports are in different SAS domains.

### Syslog Message

Disk %s is multipathed to the same SAS domain. The two disk ports must be in different SAS domains.

### Parameters

**disk** (STRING): Name of the disk.

## config.bridge.unsupp.shelf

### Severity

ALERT

### Description

This message occurs when an FC-to-SAS bridge discovers a disk shelf that it does not support.

### Corrective Action

Verify the bridge firmware is up-to-date with minimum FW rev required for support of this shelf type. If the bridge firmware supports the disk shelf but the error condition persists, remove the disk shelf.

### Syslog Message

FC-to-SAS %s bridge with S/N: %s WWN: %s has unsupported %s shelf type attached.

### Parameters

**bridgeType** (STRING): Type of FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the bridge.

**bridgeWWN** (STRING): World Wide Name of the bridge.

**shelfType** (STRING): Type of shelf.

## config.cfobothabmodule events

### config.CFObothABmodule

### Severity

ALERT

### Description

This message occurs when a disk shelf in a node is cabled for multipathing and the configuration does not support it. Multipathing means that the disk shelf is connected to one node through both disk shelf module A and module B. In HA configurations, multipathing is only supported if software disk ownership is being used and the storage shelf supports it.

### Corrective Action

For a standard HA configuration, identify the miscabled disk shelf and attach either module A or module B to the partner node. For a shared storage or multipathed HA, assign disk ownership to the partner node by using the "disk assign" command in Maintenance mode.

### Syslog Message

The disk shelf where disk %s resides is connected to this node through both disk shelf module A and module B; one of the modules must be connected to the partner node instead.

## Parameters

**disk** (STRING): Name of the disk that is cabled for multipathing.

# config.cfomixedabmodule events

## config.CFOmixedABmodule

### Severity

ALERT

### Description

This message occurs when a disk shelf attached to a node is miscabled so that one disk shelf's module A is connected to the next disk shelf's module B, or vice versa.

### Corrective Action

Identify the miscabled disk shelf, and then recable it. For FC shelves, ensure that each disk shelf's module A is connected Output to Input to the next disk shelf's module A, and that the disk shelf's module B is connected Output to Input to the next disk shelf's module B. For SAS shelves, connect the port identified with a circle on disk shelf module A to the port identified with a square on the next disk shelf's module A, and the "circle" port on module B to the next disk shelf's "square" port.

### Syslog Message

Host adapter %s has miscabled disk shelves; one disk shelf's module A is connected to the next disk shelf's module B, or vice versa.

## Parameters

**adapter** (STRING): Adapter that the miscabled disk shelf is on.

# config.differentswitchport events

## config.DifferentSwitchPort

### Severity

ALERT

### Description

This message occurs when the two ports of the disk are connected to different switch ports.

### Corrective Action

(None).

### Syslog Message

Disks on port %s:%d are also connected to port %s:%d. They need to be connected via the same port on different switches. This is an illegal configuration.

## Parameters

**switch1** (STRING): (None).

**port1** (INT): (None).

**switch2** (STRING): (None).

**port2** (INT): (None).

# config.diffpool events

## config.DiffPool

### Severity

ALERT

### Description

This message occurs when a disk is found to be multipathed to two different mirroring pools. This configuration is illegal.

### Corrective Action

See the NetApp technical support site for guidelines about assigning disks or array LUNs to SyncMirror® pools.

### Syslog Message

Disk %s is dual-attached to two separate mirrored pools. This is an illegal configuration.

### Parameters

**disk** (STRING): Name of the disk.

# config.disknoswitch events

## config.DiskNoSwitch

### Severity

ALERT

### Description

This message occurs when a disk is found to be attached to a switch, but it's not supported for a switch configuration.

### Corrective Action

Ensure that the Disk Qualification Package (DQP) is up-to-date. Visit the NetApp technical support site for information on DQP files and installation. If the error persists for more than 10 minutes after installing an updated DQP file, remove the drive.

### Syslog Message

Disk %s is not supported in a switch configuration.

### Parameters

**disk** (STRING): Name of the disk.

# config.exceededha events

## config.ExceededHA

### Severity

ERROR

## Description

This message occurs when the number of local attached FC adapters exceeds the maximum allowed.

## Corrective Action

(None).

## Syslog Message

%d Fibre Channel disk adapters were detected and the maximum allowed is only %d.

## Parameters

numHA (INT): (None).

maxHA (INT): (None).

# config.fcsameloop events

## config.FCSameLoop

### Severity

ERROR

### Description

This message occurs when a disk is multipathed to the same Fibre Channel loop.

### Corrective Action

Recable the disk shelf containing the specified disk so that it is multipathed to different Fibre Channel loops.

### Syslog Message

Disk %s is attached to the same Fibre Channel loop. The disk should be multipathed to different loops.

### Parameters

disk (STRING): Name of the disk.

# config.habadbank1 events

## config.HABadBank1

### Severity

ALERT

### Description

This message occurs when a local host adapter is connected to a different switch bank (on ports 8-15) than other adapters on this controller, which are connected to ports 0-7. This is an illegal configuration.

### Corrective Action

Connect all of the adapters of this controller to the same bank of switch ports.

### Syslog Message

Other local host adapters are connected to Switch bank 1 (ports 0-7). Host adapter %s is connected to Switch bank 2 (ports 8-15). This is an illegal configuration.

## Parameters

**ha** (STRING): Name of the host adapter.

# config.habadbank2 events

## config.HABadBank2

### Severity

ALERT

### Description

This message occurs when a local host adapter is connected to a different switch bank (on ports 0-7) than other adapters on this controller, which are connected to ports 8-15. This is an illegal configuration.

### Corrective Action

Connect all of the adapters of this controller to the same bank of switch ports.

### Syslog Message

Other local host adapters are connected to Switch bank 2 (ports 8-15). Host adapter %s is connected to Switch bank 1 (ports 0-7). This is an illegal configuration.

## Parameters

**ha** (STRING): Name of the host adapter.

# config.habadremote events

## config.HABadRemote

### Severity

ALERT

### Description

This message occurs when a remote host adapter is connected to the wrong switch bank.

### Corrective Action

Examine the cabling to determine which switch bank the local controller host adapters are attached to (the lower-numbered ports or the higher-numbered ports.) Recable the HA of the remote controller to be in the other bank of switch ports.

### Syslog Message

A remote host adapter is connected to the same switch bank as the local host adapters. This is an illegal configuration.

## Parameters

(None).

# config.invalid events

## config.invalid.PortToPort

### Severity

ERROR

### Description

This message occurs when a SAS (Serial Attached SCSI) port is attached to another SAS port. This is not a valid connection.

### Corrective Action

Check and correct the cabling; connecting two SAS ports together in this configuration is not permitted.

### Syslog Message

SAS adapter "%s" is attached to another SAS adapter.

### Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter.

## config.localswitch events

### config.localswitch

### Severity

EMERGENCY

### Description

This message occurs when a disk is found attached via both local connection and a switch. This is permitted in debug. The system will be halted in non-debug.

### Corrective Action

(None).

### Syslog Message

Disk %s is attached via a local connection and a switch. This is an illegal configuration.

### Parameters

**disk** (STRING): Name of the disk.

## config.mirroring events

### config.mirroring.failure

### Severity

ERROR

### Description

This message occurs when the name-service configuration is not replicated in the node. The out-of-sync configuration file can impact access during data access attempts.

## Corrective Action

Use the "vserver services name-service unix-user load-from-uri" or "vserver services name-service unix-group load-from-uri" command to reload the configuration file. Use the (privilege: advanced) "vserver services name-service unix-user file status" or "vserver services name-service unix-group file status" command to verify that the configuration file was successfully loaded and corresponding database was successfully generated.

## Syslog Message

Failed getting the latest %s file for Vserver %s on node %s with error %s.

## Parameters

**configuration** (STRING): Name of the configuration.

**vserver\_name** (STRING): Name of the vserver, for which the configuration was not successfully replicated.

**node\_name** (STRING): Name of the node on which configuration failed to replicate.

**error\_string** (STRING): Error encountered while extracting package.

# config.mirrornotmultipath events

## config.MirrorNotMultiPath

### Severity

ERROR

### Description

This message occurs when a disk is not multipathed but synchronous mirroring is enabled.

### Corrective Action

Multipath the disk shelf stack by also attaching it to a different host adapter via the other shelf module in each disk shelf in the stack.

### Syslog Message

Disk %s and other disks on this loop/domain are not multipathed and should be for synchronous mirroring configuration.

### Parameters

**disk** (STRING): Name of the disk.

# config.mirrorscsidisk events

## config.MirrorSCSIdisk

### Severity

ALERT

### Description

This message occurs when a SCSI disk is found in a SyncMirror® configuration.

### Corrective Action

Unconfigure SyncMirror on the aggregate that contains the SCSI disk. Configure SyncMirror using aggregates that do not contain SCSI disks.

### **Syslog Message**

Disk %s is a SCSI disk, and is illegal in a SyncMirror configuration.

### **Parameters**

**disk** (STRING): Name of the SCSI disk that was found.

## **config.mixedabmodule events**

### **config.mixedABmodule**

#### **Severity**

ERROR

#### **Description**

This message occurs when the disk shelf stack is miscabled so that one disk shelf's module A is connected to the next disk shelf's module B, or vice versa.

#### **Corrective Action**

Identify the miscabled disk shelf, and then recable it. For fibre channel shelves, ensure that each disk shelf's module A is connected Output to Input to the next disk shelf's module A, and that the disk shelf's module B is connected Output to Input to the next disk shelf's module B. For SAS shelves, connect the port identified with a circle on disk shelf module A to the port identified with a square on the next disk shelf's module A, and the "circle" port on module B to the next disk shelf's "square" port on disk shelf module B.

### **Syslog Message**

Host adapter %s has miscabled disk shelves; one disk shelf's module A is connected to the next disk shelf's module B, or vice versa. This is not a recommended configuration.

### **Parameters**

**adapter** (STRING): Name of the adapter on the miscabled disk shelf stack.

## **config.more2path events**

### **config.More2Path**

#### **Severity**

ERROR

#### **Description**

This message occurs when a disk is found to have more than two connections to the storage controller.

#### **Corrective Action**

Remove connections to the disk so there are only two.

### **Syslog Message**

Disk %s and other disks on this loop/domain are attached with more than two connections. Only two connections are supported.

### **Parameters**

**disk** (STRING): Name of the disk.



## config.move events

### config.move.fail.lost.lif.address

#### Severity

ERROR

#### Description

This message occurs when a logical interface (LIF) migration fails in such a way that there might have been a resulting internal loss of the LIF IP address.

#### Corrective Action

Use the "network interface modify" command to toggle the LIF's "-status-admin" parameter down and then up. This reinstalls the LIF address in the network stack.

#### Syslog Message

LIF migration for IP address %s from port %s to port %s failed because of an internal error: %s.

#### Parameters

**address** (STRING): LIF IPv4 or IPv6 address being moved.

**migrateFromPort** (STRING): Name of the migrate-from port.

**migrateToPort** (STRING): Name of the migrate-to port.

**internalError** (STRING): Internal error message.

## config.multipath events

### config.MultiPath

#### Severity

INFORMATIONAL

#### Description

This message occurs when the filer configuration is configured for High Availability (HA) multipathing.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

(None).

## config.noadisk events

### config.noAdisk

#### Severity

ALERT

## Description

This message occurs when the local node in an HA configuration is connected to its disk shelves through the shelf B module instead of the shelf A module. In an HA configuration, the local node must be connected to its disk shelves through the shelf A module, and the partner node must be connected to those disk shelves through the shelf B module.

## Corrective Action

For a standard HA configuration, identify and connect the local node's disk shelves to the local node using the shelf's module A. Then connect those disk shelves to the partner node using the shelf module B. For a shared storage or multipathed HA, assign disk ownership to the local and partner nodes by using the "disk assign" command in Maintenance mode.

## Syslog Message

All disk shelves are attached to the local node through shelf module B. Some of these shelves should be attached via shelf module A.

## Parameters

(None).

# config.nobconnect events

## config.noBconnect

### Severity

EMERGENCY

### Description

This message occurs when the local node in an HA configuration has no access to its partner node's disk shelves because the partner node's shelves are not connected to the local node through the shelf's module B.

### Corrective Action

For an HA configuration that uses software disk ownership, verify that both the local and partner nodes can access the disks.

### Syslog Message

The local node cannot access the partner node's disk shelves because the partner node's shelves are not connected to the local node through their shelf module B.

### Parameters

(None).

# config.nocheck events

## config.NoCheck

### Severity

NOTICE

## Description

This message occurs when the system should be halting, but it cannot because the "no-config-check?" environment variable is set to true.

## Corrective Action

(None).

## Syslog Message

This is an illegal configuration. However the Filer is not halting due to no-config-check? being set to true.

## Parameters

(None).

# config.nomultipath events

## config.NoMultiPath

### Severity

ALERT

### Description

This message occurs when a disk is found with multiple paths and it does not support them.

### Corrective Action

Remove the second path from the disk shelf.

### Syslog Message

Disk %s cannot be multipathed.

### Parameters

**disk** (STRING): Name of the disk that was found.

# config.nopartnerdisks events

## config.noPartnerDisks

### Severity

ERROR

### Description

This message occurs in a High-Availability (HA) configuration when there are no partner disks.

### Corrective Action

Ensure that both nodes of the HA pair are cabled so they have access to all the disks. Check that the adapter interface to the disk shelves is up using the "storage adapter show" command. Ensure that any switches are correctly zoned for disk visibility by the nodes.

### Syslog Message

No disks were detected for the partner; this node cannot perform takeover correctly.

**Parameters**

(None).

## **config.nopartnerluns events**

### **config.noPartnerLUNs**

**Severity**

EMERGENCY

**Description**

This message occurs in a High-Availability (HA) configuration with array LUNS when there is no attachment to both local and partner LUNs.

**Corrective Action**

Ensure that both nodes of the HA pair are cabled so they have access to all of the LUNs. Check that the adapter interface to the array is up by using the "storage adapter show" command. Ensure that any switches are correctly zoned for LUN visibility by the nodes. Verify that the array or bridge is correctly exporting its LUNs on the proper interface.

**Syslog Message**

No LUNs were detected for the partner; this node cannot perform takeover correctly.

**Parameters**

(None).

## **config.noswitchcount events**

### **config.NoSwitchCount**

**Severity**

NOTICE

**Description**

This message occurs when an attempt to obtain the switch port count fails because it does not exist.

**Corrective Action**

(None).

**Syslog Message**

Unable to obtain switch port count information for attached switches.

**Parameters**

(None).

## **config.noswitchcountfordisk events**

## **config.NoSwitchCountForDisk**

### **Severity**

NOTICE

### **Description**

This message occurs when an attempt to obtain the switch port count fails because it does not exist.

### **Corrective Action**

(None).

### **Syslog Message**

Unable to obtain switch port count information for disk %s.

### **Parameters**

**disk** (STRING): Name of the disk.

## **config.notmultipath events**

### **config.NotMultiPath**

### **Severity**

ERROR

### **Description**

This message occurs when a disk is not multipathed.

### **Corrective Action**

Multipath the disk shelf stack by also attaching it to a different host adapter via the other shelf module in each disk shelf in the stack.

### **Syslog Message**

Disk %s and other disks on this loop/domain are not multipathed and should be for improved availability.

### **Parameters**

**disk** (STRING): Name of the disk.

## **config.oneswitch events**

### **config.OneSwitch**

### **Severity**

ALERT

### **Description**

This message occurs when the disk is only connected to one switch.

### **Corrective Action**

(None).

## Syslog Message

Disk %s and other disks on this loop are only connected to one switch. For availability reasons, it should be connected to two separate switches.

## Parameters

**disk** (STRING): name of the disk

# config.pvlan events

## config.pvlan.created

### Severity

NOTICE

### Description

This message occurs when a pseudo-VLAN port is created.

### Corrective Action

(None).

### Syslog Message

Pseudo-VLAN "%s" was created on the network interface "%s".

### Parameters

**pvlanName** (STRING): Pseudo-VLAN name.

**interface** (STRING): Interface name that the pseudo-VLAN will be on.

## config.pvlan.deleted

### Severity

NOTICE

### Description

This message occurs when a pseudo-VLAN port is deleted.

### Corrective Action

(None).

### Syslog Message

Pseudo-VLAN "%s" was deleted from the network interface "%s".

### Parameters

**pvlanName** (STRING): Pseudo-VLAN name.

**interface** (STRING): Interface name that the pseudo-VLAN is on.

# config.repl events

## **config.repl.bypass.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the configuration replication bypass has been disabled. Configuration replication to any partner clusters will now resume.

### **Corrective Action**

(None).

### **Syslog Message**

The configuration replication bypass has been disabled.

### **Parameters**

(None).

## **config.repl.bypass.enabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the configuration replication bypass has been enabled. Configuration replication to any partner clusters will cease until the bypass is disabled using the 'debug crs replication-bypass disable' command.

### **Corrective Action**

(None).

### **Syslog Message**

The configuration replication bypass has been enabled.

### **Parameters**

(None).

## **config.sameha events**

### **config.sameHA**

### **Severity**

ERROR

### **Description**

This message occurs when a disk is multipathed to the same host adapter. The second disk interface should be attached to a different host adapter.

### **Corrective Action**

Move one of the connections attaching the disk shelf to a different host adapter.

### **Syslog Message**

Disk %s and other disks attached to the same port are dual-attached to the same adapter. For improved availability you should dual-attach them to separate adapters.

### **Parameters**

**disk** (STRING): Name of the disk.

## **config.sameslot events**

### **config.sameSlot**

#### **Severity**

ERROR

#### **Description**

This message occurs when a SAS disk is multipathed to host adapters in the same controller slot. The second disk interface should be attached to a host adapter in a different slot.

#### **Corrective Action**

Move one of the connections attaching the disk shelf to a host adapter in a different slot.

### **Syslog Message**

Disk %s and other disks attached to the same port are dual-attached to adapters in the same slot . For improved availability you should dual-attach them to adapters in separate slots.

### **Parameters**

**disk** (STRING): Name of the disk.

## **config.sameswitch events**

### **config.SameSwitch**

#### **Severity**

ALERT

#### **Description**

This message occurs when both ports of a disk is found to be attached to the same switch.

#### **Corrective Action**

(None).

### **Syslog Message**

Disks on port %s:%d are also connected to the same switch at port %s:%d. This is an illegal configuration.

### **Parameters**

**switch1** (STRING): name of the switch

**port1** (INT): port number of 1st connection.

**switch2** (STRING): (None).

**port2** (INT): port number of 2nd connection.



## config.sameswitchwarn events

### config.SameSwitchWarn

#### Severity

ALERT

#### Description

This message occurs when both ports of a disk are found to be attached to the same switch. Higher availability is achieved by cabling disk ports to different switches.

#### Corrective Action

Recable one of the indicated ports to a different switch.

#### Syslog Message

Disks on port %s:%d are also connected to the same switch at port %s:%d. For availability reasons, they should be connected to two separate switches.

#### Parameters

**switch1** (STRING): Name of the switch.

**port1** (INT): Port number of first connection.

**switch2** (STRING): (None).

**port2** (INT): Port number of second connection.

## config.sas2domain events

### config.SAS2Domain

#### Severity

ERROR

#### Description

This message occurs when a disk is multipathed to one SAS domain; that is, both ports of the disk are visible from the same HBA port.

#### Corrective Action

Recable the disk shelf containing the specified disk, so that the two disk ports are in different SAS domains.

#### Syslog Message

Disk %s is multipathed to the same SAS domain. The two disk ports must be in different SAS domains.

#### Parameters

**disk** (STRING): Name of the disk.

## config.scsidisk events

### config.SCSIdisk

**Severity**

NOTICE

**Description**

This message occurs when a SCSI disk is found in a CFO configuration.

**Corrective Action**

(None).

**Syslog Message**

Disk %s is a SCSI disk, and is illegal in a Clustered Failover configuration.

**Parameters**

**disk** (STRING): Name of the SCSI disk which was found.

# configbr events

## configbr.backupcompleted events

### configbr.backupCompleted

#### Severity

NOTICE

#### Description

This message occurs when a scheduled configuration backup is successfully created.

#### Corrective Action

(None).

#### Syslog Message

Scheduled configuration backup %s was created successfully.

#### Parameters

**backup** (STRING): Backup file.

## configbr.backupfailed events

### configbr.backupFailed

#### Severity

ALERT

#### Description

This message occurs when when the configuration backup cannot be created.

#### Corrective Action

Verify the root volume file system is available and has enough space to create a backup with the "volume show" command or adjust the size of the memory-backed scratch space using the "configuration backup settings modify" command (privilege: diagnostic). If the scheduled backup could not be created, create a backup file manually with the (privilege: advanced) "system configuration backup create" command, find the remote URL for uploading with the (privilege: advanced) "system configuration backup settings show -fields destination" command, and upload the backup with the (privilege: advanced) "system configuration backup upload" command.

#### Syslog Message

The %s configuration backup %s cannot be created. Error: '%s'.

#### Parameters

**type** (STRING): Type of backup, either cluster or local.

**backup** (STRING): Backup file.

**reason** (STRING): Reason for the error.

## configbr.deletefailed events

### configbr.deleteFailed

#### Severity

ERROR

#### Description

This message occurs when an old configuration backup cannot be deleted.

#### Corrective Action

Ensure that the node is accessible and healthy. If necessary, delete the backup manually with the (privilege: advanced) "system configuration backup delete" command.

#### Syslog Message

The %s configuration backup %s on node %s cannot be deleted. Error: '%s'.

#### Parameters

**type** (STRING): Type of backup, either cluster or local.

**backup** (STRING): Backup file.

**node** (STRING): Node on which the backup file exists.

**reason** (STRING): Reason for the error.

## configbr.distributefailed events

### configbr.distributeFailed

#### Severity

ERROR

#### Description

This message occurs when a configuration backup cannot be distributed to another node in the cluster.

#### Corrective Action

Ensure that the destination node is accessible and healthy. If necessary, upload the backup file to the remote URL to increase the availability of the backup file with the (privilege: advanced) "system configuration backup upload" command.

#### Syslog Message

Configuration backup %s cannot be distributed to node %s. Error: '%s'.

#### Parameters

**backup** (STRING): Backup file.

**destination\_node** (STRING): Destination node to which the backup file failed to be distributed.

**reason** (STRING): Reason for the error.

## configbr.initfailed events

## configbr.initFailed

### Severity

ERROR

### Description

This message occurs when when the configuration backup and restore feature cannot initialize the necessary software components.

### Corrective Action

Verify the root volume file system is available and has enough space to create a backup with the "volume show" command. If replicated database entries could not be populated, wait 5 minutes and verify that the backup schedules are listed with the (privilege: advanced) "system configuration backup settings show -instance" command. If the backup jobs could not be created, wait 5 minutes and verify that backup jobs are displayed for each schedule with the "job show" command.

### Syslog Message

Configuration backup and restore feature could not be initialized. Error: '%s'.

### Parameters

**reason** (STRING): Reason for the initialization failure.

## configbr.interror events

### configbr.intError

### Severity

ERROR

### Description

This message occurs when an attempt to create a configuration backup file encountered an internal software error.

### Corrective Action

Create the configuration backup file manually by using the (privilege: advanced) "system configuration backup create" command.

### Syslog Message

Configuration backup and restore detected an internal software error. Error: '%s'.

### Parameters

**error** (STRING): Internal software error.

## configbr.nodesturl events

### configbr.noDestURL

### Severity

ERROR

## Description

This message occurs when the cluster backup is saved on only one node and no offsite configuration backup destination URL is configured.

## Corrective Action

Configure an offsite configuration backup destination URL as soon as possible, using the commands "system configuration backup settings modify -destination <destination-url> -username <username-on-destination-url>" and "system configuration backup settings set-password".

## Syslog Message

Cluster backup is saved on only one node and no offsite configuration backup destination URL is configured.

## Parameters

(None).

# configbr.snapshotdelfailed events

## configbr.snapshotDelFailed

### Severity

ERROR

### Description

This message occurs when a Snapshot(tm) copy on the root volume cannot be deleted. The Snapshot copy was created during a configuration backup.

### Corrective Action

Ensure that the root volume file system is available with the "volume show" command.

### Syslog Message

Snapshot copy %s on the root volume cannot be deleted on node %s. Error: '%s'.

### Parameters

**name** (STRING): Name of the Snapshot copy.

**node** (STRING): Node on which the Snapshot copy was created.

**reason** (STRING): Reason for the error.

# configbr.unpubfailed events

## configbr.unpubFailed

### Severity

ERROR

### Description

This message occurs when the configuration backup and restore feature fails to unpublish a package by using the file replication service. The configuration backup process creates file packages and registers them with the file replication service, which stores the package in the root file system. When the configuration backup process is complete, it cleans up these files by unpublishing the package by using the file

replication service. When unpublishing fails, these files are not be cleaned up and remain in the root file system. This message does not indicate a failure in the configuration backup process, which continues despite the event. You can monitor the status of an ongoing configuration backup process by using the "job show" command. The completed configuration backup files are visible when you run the "system configuration backup show" command.

### Corrective Action

Obtain diagnostic privilege and view the file replication packages by using the "system services replication show" command. Unpublish the file replication packages by issuing the "system services replication unpublish" command, for which you also need diagnostic privilege. Check the "cluster ring show" command to ensure that the node is in quorum; management quorum is required for the unpublish command to succeed. Do not take these actions while an active configuration backup process is running.

### Syslog Message

Configuration backup and restore failed to unpublish package %s; Error: '%s'.

### Parameters

**package\_name** (STRING): Name of the package.  
**reason** (STRING): Reason for the unpublishing failure.

## configbr.uploadfailed events

### configbr.uploadFailed

#### Severity

ERROR

#### Description

This message occurs when the configuration backup cannot be uploaded to the destination URL.

#### Corrective Action

Ensure that the destination URL is reachable, the protocol in the URL is supported, and the user credentials are valid.

#### Syslog Message

Configuration backup file %s cannot be uploaded to the destination URL %s. Error: '%s'.

#### Parameters

**backup** (STRING): Backup file.  
**destination** (STRING): Destination URL.  
**reason** (STRING): Reason for the upload failure.

# connectx events

## connectx.cmdinterface events

### connectx.cmdInterface.hung

**Severity**

NOTICE

**Description**

This message occurs when the command interface of ConnectX HA interconnect hardware is temporarily hung. If the command interface fails to recover, the HA interconnect goes down.

**Corrective Action**

Typically, the command interface recovers after a few seconds. If not, an AutoSupport message is generated when the HA interconnect goes down.

**Syslog Message**

HA interconnect ConnectX hardware command interface is temporarily hung.

**Parameters**

(None).

## connectx.fatalerror events

### connectx.fatalError

**Severity**

ERROR

**Description**

This message occurs when a fatal error is polled from the ConnectX(tm) InfiniBand® adapter. Controller failover will be disabled.

**Corrective Action**

Contact NetApp technical support. ConnectX registers have been dumped to the /etc/mlxlog/ConnectX\_regdump file.

**Syslog Message**

HA interconnect HBA failure, ConnectX fatal error was polled.

**Parameters**

(None).

## connectx.ibcabledetected events

### connectx.IbCableDetected



**Severity**

INFORMATIONAL

**Description**

This message occurs when an InfiniBand cable is detected on the InfiniBand port.

**Corrective Action**

(None).

**Syslog Message**

Detected %s of length %dM on InfiniBand port %s.

**Parameters**

**cable\_type** (STRING): InfiniBand cable type.

**cable\_length** (INT): Length of the InfiniBand cable in meters.

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

## connectx.iblinkretrainnotreqd events

### connectx.IBLinkRetrainNotReqd

**Severity**

INFORMATIONAL

**Description**

This message occurs when the software detects that the previous InfiniBand link training was successful.

**Corrective Action**

(None).

**Syslog Message**

InfiniBand Link training was successful for ConnectX in slot %d, port %d.

**Parameters**

**pcie\_slot** (INT): PCIe slot in which ConnectX is present.

**port\_num** (INT): Port number where the success was detected.

## connectx.iblinkretrainreqd events

### connectx.IBLinkRetrainReqd

**Severity**

NOTICE

**Description**

This message occurs when the software detects that the previous InfiniBand link training was unsuccessful.

**Corrective Action**

(None).

## Syslog Message

InfiniBand Link training failed for ConnectX in slot %d, port %d. Link Retraining Request status %d.

## Parameters

**pcie\_slot** (INT): PCIe slot in which ConnectX is present.  
**port\_num** (INT): Port number where the issue was detected.  
**status** (INT): Status of the link retraining request.

# connectx.ibqsfpdumpctrl events

## connectx.IbQsfpDumpCtrl

### Severity

ERROR

### Description

This message occurs when InfiniBand retimer programming fails due to a quad small form-factor pluggable (QSFP) transceiver register dump failure, causing the control registers of QSFP to be dumped.

### Corrective Action

Automatic retry of InfiniBand retimer programming will be done and an AutoSupport(tm) message will be generated if two retries fail.

## Syslog Message

InfiniBand retimer programming failed on port %s due to %s. Dumping registers: control 0x%0x, data 0x%0x, timeout 0x%0x, clock 0x%0x.

## Parameters

**port\_str** (STRING): InfiniBand port in which the cable is plugged.  
**reason** (STRING): Reason for the failure of InfiniBand retimer programming.  
**ctrlreg** (INTHEX): Contents of the QSFP control register in hexadecimal notation.  
**datareg** (INTHEX): Contents of the QSFP data register in hexadecimal notation.  
**timeoutreg** (INTHEX): Contents of the QSFP timeout register in hexadecimal notation.  
**clockreg** (INTHEX): Contents of the QSFP clock register in hexadecimal notation.

# connectx.ibqsfpdumpdata events

## connectx.IbQsfpDumpData

### Severity

ERROR

### Description

This message occurs when InfiniBand retimer programming fails due to a quad small form-factor pluggable (QSFP) transceiver parse error, and 128 bytes of QSFP data are dumped.

### Corrective Action

Automatic retry of InfiniBand retimer programming will be done and an AutoSupport(tm) message will be generated if two retries fail.

## Syslog Message

InfiniBand retimer programming failed on port %s due to %s. Dumping registers: %s.

## Parameters

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

**reason** (STRING): Reason for the failure of InfiniBand retimer programming.

**qsfp\_data** (STRING): Contents of QSFP data registers from offset 128 to 256 in hexadecimal notation.

# connectx.ibretimerfailcount events

## connectx.IbRetimerFailCount

### Severity

ERROR

### Description

This message occurs when InfiniBand retimer programming fails due to one of the following reasons: InfiniBand device not found, GPIO read failure or retimer chip failure.

### Corrective Action

Automatic retry of InfiniBand retimer programming will be done and an AutoSupport(tm) message will be generated if two retries fail.

## Syslog Message

InfiniBand retimer programming failed on port %s due to %s.

## Parameters

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

**reason** (STRING): Reason for the failure of InfiniBand retimer programming.

# connectx.ibretimerprogrmpass events

## connectx.IbRetimerProgrmPass

### Severity

INFORMATIONAL

### Description

This message occurs when InfiniBand retimer programming is successfully completed after the system detects that a cable is plugged in.

### Corrective Action

(None).

## Syslog Message

InfiniBand retimer programming was successful on port %s.

## Parameters

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

## connectx.ibretimerskipprogrm events

### connectx.IbRetimerSkipProgrm

#### Severity

INFORMATIONAL

#### Description

This message occurs when InfiniBand retimer programming is skipped after the system detects that a cable is plugged in, because there was no change in the cable's type and length.

#### Corrective Action

(None).

#### Syslog Message

InfiniBand retimer programming was skipped on port %s for %s cable.

#### Parameters

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

**cable\_type** (STRING): InfiniBand cable type.

## connectx.ibunsupportcable events

### connectx.IbUnsupportCable

#### Severity

ERROR

#### Description

This message occurs when an unsupported InfiniBand cable is detected on the InfiniBand port.

#### Corrective Action

Unplug the unsupported InfiniBand cable and plug in a supported cable. Search the support site "Hardware Universe" for information about supported InfiniBand cables.

#### Syslog Message

Detected unsupported %s on InfiniBand port %s.

#### Parameters

**cable\_type** (STRING): InfiniBand cable type.

**port\_str** (STRING): InfiniBand port in which the cable is plugged.

## connectx.portdisabled events

### connectx.portDisabled

#### Severity

ALERT

**Description**

This message occurs when the HA interconnect port experiences persistent node advertisement send timeouts. Port reinitialization does not correct the problem. The port is disabled.

**Corrective Action**

Reboot the controller. Contact NetApp technical support if the problem persists.

**Syslog Message**

Node advertisement send timed out for the ConnectX device in slot %d, port %d. The port is disabled.  
Device registers have been dumped to the /etc/mlxlog/ConnectX\_regdump file.

**Parameters**

**phys\_slot** (INT): Physical slot in which ConnectX is present.

**port\_num** (INT): Port number that is disabled.

# copy events

## copy.dst events

### copy.dst.raid.convert.fail

#### Severity

NOTICE

#### Description

This event is issued when a SnapMirror or aggregate copy operation attempts to update a flexible volume RAID block on the destination but is unable to do so.

#### Corrective Action

Verify the source and destination have flexible volumes with compatible versions.

#### Syslog Message

Unable to convert volume RAID block %llu during copy transfer.

#### Parameters

**block** (LONGINT): Block with the error

# coredump events

## coredump.all events

### coredump.all.disks.used

**Severity**

ERROR

**Description**

All disks available to coredump already contain unsaved cores.

**Corrective Action**

There are three choices. First, add more disks to the system. Second, delete any unsaved cores by running "savecore -k". Third, save any unsaved cores by running "savecore". It may be necessary to run "savecore -f" to save the cores, if previous attempts to save these cores have failed.

**Syslog Message**

No disks available for dumpcore

**Parameters**

(None).

## coredump.compress events

### coredump.compress.failed

**Severity**

ERROR

**Description**

An error was encountered while compressing data for the coredump. This core will be deleted.

**Corrective Action**

(None).

**Syslog Message**

Error compressing core data: %s

**Parameters**

msg (STRING): The error message

## coredump.compression events

### coredump.compression.details

**Severity**

INFORMATIONAL

## Description

This message occurs when a compressed kernel core is generated. Details about compression are reported.

## Corrective Action

(None).

## Syslog Message

%s, compression time %llu msec, compression count %llu, bytes before compression %llu, bytes after compression %llu, compression ratio %llu%%, compression bandwidth %llu MB/s, total compression time %llu msec, all-zero chunk count %llu.

## Parameters

**type** (STRING): Type of dump. Possible types are compressed core, compressed spraycore, spraycore and sparecore.

**compressTime** (LONGINT): Compression time, in milliseconds.

**compressCount** (LONGINT): Number of compressions.

**compressBytesIn** (LONGINT): Number of bytes before compression.

**compressBytesOut** (LONGINT): Number of bytes after compression.

**compressRatio** (LONGINT): Compression ratio expressed as a percentage.

**compressBandwidth** (LONGINT): Compression bandwidth, in MBps.

**totalCompressTime** (LONGINT): Total compression time, in milliseconds.

**allZeroChunks** (LONGINT): Number of all-zero chunks.

# coredump.coreid events

## coredump.coreid.not.found

### Severity

INFORMATIONAL

### Description

Could not find the specified core

### Corrective Action

Double-check the core id, and try again.

### Syslog Message

Core %d could not be found

### Parameters

**coreid** (INT): The core id that could not be found

# coredump.delete events

## coredump.delete.denied

### Severity

INFORMATIONAL



## Description

This message occurs when a requested coredump delete cannot occur because an aggregate relocation or a takeover is in progress. The system cannot delete a core while either an aggregate relocation or a takeover is in progress.

## Corrective Action

Retry the coredump delete after aggregate relocation or takeover completes.

## Syslog Message

Coredump delete disallowed because an aggregate relocation or a takeover in progress.

## Parameters

(None).

# coredump.disk events

## coredump.disk.write.details

### Severity

INFORMATIONAL

### Description

This message occurs when a kernel core is written to disk. Details about disk operations are reported.

### Corrective Action

(None).

### Syslog Message

%s, disk write time %llu msec, disk write count %llu, disk write bytes %llu, disk write bandwidth %llu MB/s.

### Parameters

**type** (STRING): Type of dump. Possible types are compressed core, compressed spraycore, spraycore and sparecore.

**writeTime** (LONGINT): Disk write time, in milliseconds.

**writeCount** (LONGINT): Number of disk writes.

**writeBytes** (LONGINT): Number of bytes written to disk.

**writeBandwidth** (LONGINT): Disk write bandwidth, in MBps.

# coredump.dr events

## coredump.dr.dump.not.started

### Severity

ERROR

### Description

This message occurs when coredump subsystem detects that disaster recovery (DR) partner sparecore dump did not start.

### Corrective Action

Log in to service processor and power cycle the DR partner that may not have started dumping core.

### Syslog Message

DR partner sparecore dump failed to start.

### Parameters

(None).

## coredump.dr.spare.abort

### Severity

INFORMATIONAL

### Description

This message occurs when coredump subsystem attempts to abort any sparecore dump that might be in progress on disaster recovery (DR) partner. This process can take up to 60 seconds. During this time, other coredump operations are not serviced.

### Corrective Action

(None).

### Syslog Message

Pausing to abort DR partner sparecore dump on %s because %s.

### Parameters

**diskuid** (STRING): UID of DR partner sparecore dump disk.

**reason** (STRING): Reason for aborting the DR partner sparecore dump.

## coredump.dr.spare.noAbort

### Severity

ERROR

### Description

This message occurs when coredump subsystem is unable to abort any sparecore dump that might be in progress on disaster recovery (DR) partner.

### Corrective Action

Log in to service processor and power cycle the DR partner that is still dumping core.

### Syslog Message

DR partner sparecore dump could not be aborted on disk %s because %s.

### Parameters

**diskuid** (STRING): UID of DR partner sparecore dump disk.

**reason** (STRING): Reason for abort failure.

# coredump.dump events

## coredump.dump.abort

### Severity

ERROR

### Description

This message occurs when there is an unsaved core on the dump device. The abort prevents a new crash from overwriting the original dump core, which is expected to contain more interesting information for analysis.

### Corrective Action

If the default behavior is not desired, then set the boot option 'save-latest-core?' to true.

### Syslog Message

Dump is aborted because a core in the dump device is not saved yet.

### Parameters

(None).

## coredump.dump.disk.failed

### Severity

NOTICE

### Description

This message occurs once for each core disk that has read or write errors while dumping the core. Each disk is marked as failed, so further writes to the disks are dropped. Dumpcore tries to recover from this error, and continue dumping the core.

### Corrective Action

(None).

### Syslog Message

Disk %s had %u read and %u write error(s) during dumpcore.

### Parameters

**disk** (STRING): Name of this core disk marked failed with read or write errors.

**readErrorCount** (INT): Number of read errors seen on this core disk.

**writeErrorCount** (INT): Number of write errors seen on this core disk.

## coredump.dump.disk\_failed.count

### Severity

ERROR

### Description

During the last dumpcore attempt, disk write errors are ignored, and dumpcore continues. In this case, too many disks failed during the last dump attempt for savecore to create a core file. The core will be deleted, and no core will be available from this panic.

## Corrective Action

(None).

## Syslog Message

%d of %d disks failed during dumpcore

## Parameters

**failed** (INT): Number of failed disks

**total** (INT): Number of disks used

## coredump.dump.disks.count

### Severity

ERROR

### Description

This message occurs when the system fails to write a coredump due to lack of disk space.

### Corrective Action

Disk space needed for a coredump can be reduced by enabling sparse cores, which is the default setting. Verify that sparse cores are enabled by using the "system node coredump config show" command. If not, enable them by using the "system node coredump modify -node local -sparse-enabled true" command. If there is still not enough disk space to write a coredump, add a spare disk that is at least 0.6 times the sum of primary memory size and NVRAM memory size. You can obtain sizes of primary memory and NVRAM memory by using the "run local sysconfig" command and reading entries for system board and NVRAM.

## Syslog Message

Not enough disks to write core dump.

## Parameters

(None).

## coredump.dump.dont

### Severity

INFORMATIONAL

### Description

This message occurs when a core dump is not performed because the system was configured to not dump cores or dump is disabled in the code path that caused the system to panic. Any of the following configuration methods can disable coredump: - Environment variable "dont-dump-core?" is "true" - Parameter "-coredump-attempts" in the command "system node coredump config" is "0".

### Corrective Action

If the reason is "dump disabled in the current code path", there is no corrective action that can be taken. If the reason is "coredump configured to not dump cores", 1. "-coredump-attempts" parameter might be set to "0". Check the "Max Dump Attempts" value in the output of the "system node coredump config show -node node\_name" command. If "0", use the "system node coredump config modify ←node node\_name> -coredump-attempts <number\_of\_attempts>" command to set a non-zero value. The default is 2 attempts. 2. "dont-dump-core?" environment variable might be set to "true". From the nodeshell, use the (priv: diag) "bootargs dump" command to check whether "dont-dump-core?" exists and is set to "true". To unset the variable, use the "bootargs unset dont-dump-core?" command.

**Syslog Message**

Coredump not performed because %s.

**Parameters**

**reason** (STRING): Reason that coredump was not performed.

**coredump.dump.failed****Severity**

EMERGENCY

**Description**

Coredump failed to complete.

**Corrective Action**

Previously logged coredump.dump.\* events should help explain the exact cause of the failure.

**Syslog Message**

Dumpcore failed

**Parameters**

(None).

**coredump.dump.find.all.used****Severity**

ERROR

**Description**

Unsaved cores occupy all the disks available to dumpcore.

**Corrective Action**

Save the unsaved cores using "savecore" or "savecore -f". If the cores aren't needed, run "savecore -k" to delete them.

**Syslog Message**

All available disks contain unsaved cores

**Parameters**

(None).

**coredump.dump.find.disks.failed****Severity**

ERROR

**Description**

Could not find any suitable disks for dumping a core.

**Corrective Action**

Make sure there are non-broken disks connected to (and owned by) the panicing host.

**Syslog Message**

No valid disks found

**Parameters**

(None).

**coredump.dump.find.spare.failed****Severity**

ERROR

**Description**

A sparecore dump is supposed to be done, but no suitable spare disk can be found. In these situations, it is more important to allow a fast takeover than it is to get the core, so the dump is terminated, and takeover begins.

**Corrective Action**

Add a spare disk to the system, or disable the "cf.takeover.on\_panic" option.

**Syslog Message**

No valid spare disk found

**Parameters**

(None).

**coredump.dump.nvram.copy.failed****Severity**

ERROR

**Description**

Error while copying NVRAM into main memory during dumpcore. The resulting coredump will have dummy data in the NVRAM segment.

**Corrective Action**

(None).

**Syslog Message**

Error while copying NVRAM contents into main memory

**Parameters**

(None).

**coredump.dump.raid.notready****Severity**

ERROR

**Description**

RAID has not yet processed the labels on the disks, so dumpcore cannot proceed. This is typically the case for panics very early during boot.

**Corrective Action**

(None).

**Syslog Message**

RAID not ready to dump cores

**Parameters**

(None).

**coredump.dump.recursive.late****Severity**

ERROR

**Description**

A recursive panic situation has been encountered. The location of the second panic makes it unsafe to attempt another dump.

**Corrective Action**

(None).

**Syslog Message**

Late recursive panic caused dumpcore to abort

**Parameters**

(None).

**coredump.dump.save.latest.core****Severity**

ERROR

**Description**

The boot option `save-latest-core?` is set to true. This will discard all unsaved cores during coredump.

**Corrective Action**

If the behavior is not desired then set the boot option `save-latest-core?` to false.

**Syslog Message**

Saving the latest core and discarding all unsaved cores.

**Parameters**

(None).

**coredump.dump.scrub.skipped**

**Severity**

NOTICE

**Description**

This message occurs when the coredump process was asked to scrub sensitive memory areas but could not do so. The dump must be aborted to avoid saving such memory areas to disk.

**Corrective Action**

(None).

**Syslog Message**

Could not scrub protected memory regions.

**Parameters**

(None).

**coredump.dump.segments.count****Severity**

ERROR

**Description**

Dumpcore ran out of space for the core segments in the header. This core cannot be dumped.

**Corrective Action**

Disable the "coredump.metadata\_only" option.

**Syslog Message**

Out of segment space in the core header

**Parameters**

(None).

**coredump.dump.spare.no.space****Severity**

ERROR

**Description**

The selected spare disk might not be big enough to hold the core. The dump attempt has been aborted.

**Corrective Action**

Make sure a spare disk is available that is at least as big as the sum of main memory and NVRAM. Set the option "coredump.metadata\_only" to "on". If the problem persists, turn off the option "cf.takeover.on\_panic". This will disable the sparecore feature, and delay takeovers until the panicing partner has finished dumping the core.

**Syslog Message**

Spare disk too small for coredump



## Parameters

(None).

## coredump.dump.started

### Severity

INFORMATIONAL

### Description

Basic info about start of dumpcore

### Corrective Action

(None).

### Syslog Message

%s starting with %d disks

### Parameters

**type** (STRING): Type of dump. Possible types are compressed, sprayed, and sparecore.

**disks** (INT): Number of disks being used

## coredump.dump.time

### Severity

INFORMATIONAL

### Description

Information about how long dumpcore took to complete.

### Corrective Action

(None).

### Syslog Message

%s completed in %d seconds

### Parameters

**type** (STRING): Type of dump. Possible types are compressed, sprayed, and sparecore.

**seconds** (INT): Number of seconds spent on the dump

## coredump.dump.write.blocks.failed

### Severity

ERROR

### Description

Failed to write dump data blocks to disk. If the write failed because of disk problems, dumpcore will restart, without the disk that failed. If the write failed for other reasons (eg. attempt to write to a bad position on the disk), dumpcore will be aborted.

### **Corrective Action**

(None).

### **Syslog Message**

Error dumping %llu blocks starting at block %llu to disk %s

### **Parameters**

**blocks** (LONGINT): Number of blocks trying to be dumped

**start** (LONGINT): Starting block number

**disk** (STRING): Name of disk on which the write failed

## **coredump.dump.write.error**

### **Severity**

ERROR

### **Description**

An i/o error occurred while attempting to write a coredump to disk.

### **Corrective Action**

(None).

### **Syslog Message**

An i/o error occurred while attempting to dump core.

### **Parameters**

(None).

## **coredump.dump.write.header.failed**

### **Severity**

ERROR

### **Description**

Failed to write the core header to disk during dumpcore.

### **Corrective Action**

(None).

### **Syslog Message**

Error writing core header to block %llu on disk %s

### **Parameters**

**header** (LONGINT): Block number of the core header block

**disk** (STRING): Name of disk on which the write failed

## **coredump.encryption events**

## **coredump.encryption.disabled**

### **Severity**

ERROR

### **Description**

This message occurs when encryption is configured on the controller, but the coredump subsystem cannot encrypt data. A system panic in this situation results in no core file being produced.

### **Corrective Action**

(Call support).

### **Syslog Message**

Coredump encryption cannot encrypt data. Core file will not be saved in this state.

### **Parameters**

(None).

## **coredump.encryption.restored**

### **Severity**

NOTICE

### **Description**

This message occurs when encryption is configured on the controller, the coredump subsystem experienced a transient problem encrypting data, but has since recovered and will be able to save encrypted core dump files in the event of a system panic.

### **Corrective Action**

(None).

### **Syslog Message**

Coredump encryption experienced a failure in the past but has recovered. Core files can now be saved.

### **Parameters**

(None).

## **coredump.findcore events**

### **coredump.findcore.error**

### **Severity**

ERROR

### **Description**

This message occurs when the core dump process detects an error while checking for an unsaved core in the dump device. The dump is not aborted.

### **Corrective Action**

(None).

## Syslog Message

Ignored the error (%d: %s) when trying to find an unsaved core in the dump device.

## Parameters

**errCode** (INT): Error code.

**errString** (STRING): Error code represented as a string.

## coredump.findcore.nospace

### Severity

ERROR

### Description

There are too many unsaved cores on the system to handle them all at one time. Some cores are being ignored.

### Corrective Action

Run "savecore" multiple times to save all the cores. Alternatively, run "savecore -k" to delete all of them.

## Syslog Message

Too many unsaved cores

## Parameters

(None).

## coredump.findcore.partial

### Severity

ERROR

### Description

This message occurs during a core dump, when only a partial core is found. Although enough of this core is present for the core to be partially saved, there might be errors while doing so. These occur when data essential to the core creation process is lost with the missing disks.

### Corrective Action

Run the "savecore" command to attempt to save the partial core, or run the "savecore -k" command to delete all unsaved cores.

## Syslog Message

Partial core %s is missing %d of %d disks.

## Parameters

**corefile** (STRING): Name of the kernel core file.

**missing** (INT): Number of disks that are missing.

**total** (INT): Total number of disks used in this dump.

## coredump.findcore.partial.nosave

**Severity**

ERROR

**Description**

This message occurs when a partial core is found with too much missing to be saved.

**Corrective Action**

Run the "savecore -k" command to delete all unsaved cores.

**Syslog Message**

Cannot save partial core %s, missing %d of %d disks.

**Parameters**

**corefile** (STRING): Name of the kernel core file.

**missing** (INT): Number of disks that are missing.

**total** (INT): Total number of disks used in this dump.

## coredump.host events

### coredump.host.spare.none

**Severity**

INFORMATIONAL

**Description**

This message occurs when no sparecore disk is found at boot/takeover time.

**Corrective Action**

(None).

**Syslog Message**

No sparecore disk was found for host %d.

**Parameters**

**host** (INT): Host for which no sparecore disk was found.

### coredump.host.spare.save

**Severity**

INFORMATIONAL

**Description**

This message occurs to tell you which disk was identified as the sparecore disk at boot/takeover time.

**Corrective Action**

(None).

**Syslog Message**

Disk %s was identified as the sparecore disk at boot/takeover time for host %d.

## Parameters

**diskname** (STRING): Name of the sparecore disk.  
**host** (INT): Host that the sparecore disk belongs to.

## coredump.host.spare.set

### Severity

INFORMATIONAL

### Description

This message occurs when a disk is identified as a sparecore disk and the host's sparecoreinfo is updated during assimilation.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**disk** (STRING): Name of the sparecore disk  
**host** (INT): Host whose sparecoreinfo was updated.

## coredump.job events

### coredump.job.notstarted

#### Severity

ERROR

#### Description

This message occurs when the system can not automatically save a core failed.

#### Corrective Action

Use the 'system node coredump save' command to save the core.

#### Syslog Message

(None).

#### Parameters

(None).

### coredump.job.notswitchedback

#### Severity

ERROR

#### Description

This message occurs when the system cannot automatically save cores because a MetroCluster(tm) switchback was not performed.

## Corrective Action

Return storage to the MetroCluster member by using the "metrocluster switchback" command, and then use the "coredump save-all" command to save any unsaved cores.

## Syslog Message

System cannot automatically save cores because a MetroCluster switchback was not performed.

## Parameters

(None).

# coredump.long events

## coredump.long.queued.msg

### Severity

NOTICE

### Description

This message occurs when the coredump manager thread detects that a message has been queued longer than expected, as defined by an internal threshold. System features and performance are not affected.

### Corrective Action

(None).

### Syslog Message

Coredump manager starts processing a long queued message: "%s" queued "%lu" ms.

### Parameters

**message\_type** (STRING): Coredump message type.  
**queue\_time** (LONGINT): Message queued duration, in milliseconds.  
**threshold** (LONGINT): Message queued threshold, in milliseconds.

## coredump.long.running.msg

### Severity

NOTICE

### Description

This message occurs when the coredump manager thread has finished processing a message that took longer than expected, as defined by an internal threshold.

### Corrective Action

(None).

### Syslog Message

Coredump manager finished processing a long running message: %s took %lu ms.

### Parameters

**message\_type** (STRING): Coredump message type.  
**duration** (LONGINT): Message processing duration, in milliseconds.  
**threshold** (LONGINT): Message processing threshold, in milliseconds.

# coredump.mcc events

## coredump.mcc.nso.veto

### Severity

INFORMATIONAL

### Description

This message occurs when the coredump subsystem vetoes MetroCluster(tm) negotiated switchover because a coredump save is in progress.

### Corrective Action

Use the "node coredump status" command to verify the completion of the active coredump save, or use the "-override-vetoes true" option when performing MetroCluster switchover. Using the "-override-vetoes true" option will abort the active coredump save.

### Syslog Message

MetroCluster switchover vetoed because a coredump save was in progress.

### Parameters

(None).

# coredump.micro events

## coredump.micro.completed

### Severity

INFORMATIONAL

### Description

Microcore generation has completed.

### Corrective Action

(None).

### Syslog Message

Microcore (%s) generation completed

### Parameters

**file** (STRING): Pathname of microcore that was created

## coredump.micro.failed

### Severity

ERROR

### Description

Microcore generation has failed.



**Corrective Action**

(Call support).

**Syslog Message**

Generation of microcore failed

**Parameters**

(None).

## coredump.nosavecore events

**coredump.nosavecore.noflush****Severity**

NOTICE

**Description**

This message occurs when the process that saves a coredump is not able to flush to the boot device at the end of coredump.

**Corrective Action**

(None).

**Syslog Message**

Coredump was not able to flush the remaining data to the boot device. Error: %llu.

**Parameters**

`dumpcore_error` (LONGINTEX): Coredump error code.

## coredump.nvlog events

**coredump.nvlog.replay.dump.dont****Severity**

NOTICE

**Description**

This message occurs when one or more unsuccessful attempts have been made to process information in the WAFL® NVlog after a previous failure. The number of allowed core dumps has been exceeded and this core dump will not be taken.

**Corrective Action**

The node will reboot. No further action is necessary.

**Syslog Message**

Core dump was not taken for the latest failure while processing the WAFL NVlog.

**Parameters**

(None).

# coredump.save events

## coredump.save.aborted

### Severity

NOTICE

### Description

This message occurs when the savecore operation is aborted while saving a kernel core. The core is preserved for a future iteration of savecore.

### Corrective Action

(None).

### Syslog Message

%s processing savecore aborted%s%s.

### Parameters

**file** (STRING): Name of the kernel core file being created.

**operation** (STRING): Type of operation that is being attempted when the abort occurs.

**action** (STRING): Type of action being taken when the abort occurs.

## coredump.save.attempts.count

### Severity

NOTICE

### Description

Previous attempts to save this core have failed, either due to reboots or panics. To prevent a panic loop, savecore is aborting this attempt.

### Corrective Action

Use either "savecore -f" to save the core, or "savecore -k" to erase the core.

### Syslog Message

Too many attempts to save this core

### Parameters

(None).

## coredump.save.completed

### Severity

NOTICE

### Description

This message occurs when all the core segments belonging to a coredump have been saved.

### Corrective Action

(None).

## Syslog Message

Saved core %s for node %s in %llu seconds. Core size is %llu bytes.

## Parameters

**file** (STRING): Primary core segment file that was created.

**node** (STRING): Name of the node whose core was saved.

**seconds** (LONGINT): Number of seconds spent on saving all the core segments.

**size** (LONGINT): Size of primary core segment in bytes.

**chksum** (STRING): Checksum of primary core segment file's data.

## coredump.save.denied

### Severity

INFORMATIONAL

### Description

This message occurs when the coredump subsystem cannot finish saving a core because an aggregate relocation or a takeover is in progress. The system cannot save a core while either an aggregate relocation or a takeover is in progress.

### Corrective Action

(None).

## Syslog Message

Coredump save disallowed because an aggregate relocation or a takeover in progress.

## Parameters

(None).

## coredump.save.disk.missing

### Severity

INFORMATIONAL

### Description

A disk being used by coredump is now missing. This could be due to the ownership of the disk changing, the disk being removed, etc. Most savecore operations will proceed, but if further errors are encountered, or an expected core cannot be found, corrective action will be necessary.

### Corrective Action

If this disk has failed, no action is possible. If this message was seen while running a "savecore -s", "savecore -i", or "savecore -l" operation, replace the missing disk. This might be done by undoing the ownership change, or reinserting the missing disk. If this message was seen while saving the core, and no further errors were seen, no action is necessary. If more errors were seen, replace the missing disk as above, and run "savecore" again.

## Syslog Message

Disk %s being used by coredump is missing

## Parameters

**disk** (STRING): Name of the missing disk

## **coredump.save.diskio.zone**

### **Severity**

ERROR

### **Description**

A disk IO operation was attempted outside of the allowed regions on the disk. This behavior indicates a corrupt core, and it will be deleted.

### **Corrective Action**

(None).

### **Syslog Message**

IO to %d block(s) starting at %llu on disk %s is not allowed

### **Parameters**

**count** (INT): Number of blocks in IO request

**start** (LONGINT): Starting block number on the disk

**disk** (STRING): Name of the disk causing trouble

## **coredump.save.error**

### **Severity**

NOTICE

### **Description**

A correctable error was encountered while saving the core. Savecore will continue saving other cores.

### **Corrective Action**

Once savecore finishes, address errors that may have appeared, and restart savecore.

### **Syslog Message**

%s processing encountered error

### **Parameters**

**file** (STRING): The core file that savecore was attempting to create

## **coredump.save.failed**

### **Severity**

NOTICE

### **Description**

An unrecoverable error was encountered while saving the core. This core will be deleted, and savecore will continue saving other cores.

### **Corrective Action**

(None).

### **Syslog Message**

%s processing failed

## Parameters

**file** (STRING): The core file that savecore was attempting to create

## coredump.save.internal.error

### Severity

ERROR

### Description

An internal coredump error was encountered. This indicates the dumped core is corrupt, and will be deleted.

### Corrective Action

(None).

### Syslog Message

Internal coredump error

### Parameters

(None).

## coredump.save.nospace

### Severity

ALERT

### Description

This message occurs when there is not enough space in the root volume to save the core. Savecore is making a worst-case estimate when determining how much space might be needed to save the core, and does not start unless that much space is available.

### Corrective Action

Try performing the following steps to make sufficient room in root volume: Delete old or unneeded core files by using the "coredump delete" command. Delete old or unneeded core file segments by using the "coredump segment delete" command. Delete any Snapshot@ copies that include the deleted files. Adjust minimum free space on the root file system by using "coredump config modify" command. Once there is sufficient room, take the following steps to save the cores: To save all core files using "coredump save-all" command, run the "coredump status -instance" command and ensure that "Space Available On Internal Filesystem" is greater than "Space Needed to Save All Unsaved Cores" plus "Minimum Free Bytes on Root Filesystem". To save a specific core file, run the "coredump save" command, when available space of "Space Needed To Save Core" (from the "coredump show" command) plus "Minimum Free Bytes on Root Filesystem" is sufficient.

### Syslog Message

Available space (%llu %dKB blocks) is not enough to save the core image (requires up to %llu blocks) and maintain the minimum free space of %llu blocks.

### Parameters

**avail** (LONGINT): Number of blocks available for saving the core image.

**blocksize** (INT): Block size in kilobytes.

**needed** (LONGINT): Number of blocks needed for saving the core image.

**minfree** (LONGINT): Number of blocks reserved for normal system operation.

## **coredump.save.partial.not.ok**

### **Severity**

INFORMATIONAL

### **Description**

Can not save partial core that was dumped using disks with different sizes. Savecore can not determine the size of the missing disks. Attempting a save in this condition will most likely end up in an error.

### **Corrective Action**

Run "savecore -k" to delete all unsaved cores.

### **Syslog Message**

Cannot save partial core (%s) that used disks with different sizes, missing %llu of %llu disks

### **Parameters**

**panic** (STRING): Panic string of the partial core.

**missing** (LONGINT): The number of disks that are missing

**total** (LONGINT): Total number of disks used in this dump

## **coredump.save.read.header.failed**

### **Severity**

NOTICE

### **Description**

An error was encountered while trying to read the core header block from disk.

### **Corrective Action**

(None).

### **Syslog Message**

Could not read core header block (%llu) from disk %s

### **Parameters**

**block** (LONGINT): Block number being searched for the core header

**disk** (STRING): Name of the disk that could not be read

## **coredump.save.started**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the savecore operation starts saving a kernel core.

### **Corrective Action**

(None).

### **Syslog Message**

Saving %llu MB to %s via %s ("%s") for node %s.

## Parameters

**size** (LONGINT): Number of MB being saved.  
**file** (STRING): Name of the kernel core file to be created.  
**dumpcoretype** (STRING): Kernel coredump type.  
**panic** (STRING): Panic string of the kernel core being saved.  
**node** (STRING): Name of the node whose core is being saved.

# coredump.segment.events

## coredump.segment.auto.fail

### Severity

NOTICE

### Description

This message occurs when the system encounters an error while trying to send a core segment request.

### Corrective Action

(None).

### Syslog Message

Failed to start automatic segmenting for core file, %s , for host %d.

## Parameters

**file** (STRING): Name of the core file that would have been segmented.  
**host** (INT): Host that owns the core file.

## coredump.segment.auto.pfail

### Severity

INFORMATIONAL

### Description

This message occurs when the system cannot start the segmenting of a core file because the node is no longer in takeover mode.

### Corrective Action

(None).

### Syslog Message

Cannot start segmenting a core file, %s. Node is no longer in takeover mode.

## Parameters

**file** (STRING): Name of core file that would have been segmented.

## coredump.segment.completed

### Severity

NOTICE

## Description

This message occurs when the segmenting of a full core file has been completed.

## Corrective Action

(None).

## Syslog Message

Saved core segment %s for node %s. Core size is %llu bytes and core data checksum is %s.

## Parameters

**file** (STRING): Name of the core file segmented.

**node** (STRING): Name of the node whose core segment was saved.

**size** (LONGINT): Size of core segment in bytes.

**chksum** (STRING): Checksum of the core segment file's data.

# coredump.shutdown events

## coredump.shutdown.trace

### Severity

NOTICE

### Description

This message occurs when a core is about to be generated. Durations of various activities at that time are reported.

### Corrective Action

(None).

### Syslog Message

coredump shutdown times: kmod dumper start %llu, panic\_info set %llu, sf\_dumpcore sent %llu, sparecore id sent %llu, fmot\_outage info sent %llu, disk\_dump\_start done %llu, coredump begin %llu (all in msec).

### Parameters

**kmodDumperStartTime** (LONGINT): Duration from panic start to dump processing start, in milliseconds.

**panicInfoTime** (LONGINT): Duration until panic information is collected, in milliseconds.

**sfDumpcoreSentTime** (LONGINT): Duration until SF\_DUMPCORE state is sent over the interconnect, in milliseconds.

**sparecoreSentTime** (LONGINT): Duration for sending the sparecore ID over the interconnect, in milliseconds.

**fmotOutageSentTime** (LONGINT): Duration for sending FMOT (Failover Monitor Outage Tracker) Outage information over the interconnect, in milliseconds.

**diskDumpStartTime** (LONGINT): Duration until disk dump start stage, in milliseconds.

**coredumpBeginTime** (LONGINT): Duration until the start of the actual core dump, in milliseconds.

# coredump.spare events

## coredump.spare.abort



**Severity**

INFORMATIONAL

**Description**

This message occurs when coredump attempts to terminate any sparecore dump that might be in progress on the partner. This process can take up to 60 seconds. During this time, other coredump operations are not serviced.

**Corrective Action**

(None).

**Syslog Message**

Pausing to abort partner sparecore dump because %s.

**Parameters**

**reason** (STRING): Reason for aborting the partner sparecore dump.

**coredump.spare.abort.failed****Severity**

ERROR

**Description**

Unable to reserve the sparecore disk. This disk cannot be used by RAID until it is successfully reserved.

**Corrective Action**

If the other partner can be booted, do a giveback. Otherwise, periodic attempts will automatically be made to reserve the disk in question.

**Syslog Message**

Sparecore disk could not be reserved

**Parameters**

(None).

**coredump.spare.abort.release****Severity**

NOTICE

**Description**

This message occurs when the system encounters an error while trying to release the disk reservation from a disk.

**Corrective Action**

(None).

**Syslog Message**

Could not release disk reservation from disk %s (0x%x)

## Parameters

**diskname** (STRING): Name of the disk whose disk reservation could not be released.

**diskid** (INHEX): ID of the disk whose disk reservation could not be released.

## coredump.spare.dump

### Severity

INFORMATIONAL

### Description

This message occurs to tell you which disk was used as the sparecore disk at dumpcore time.

### Corrective Action

(None).

### Syslog Message

Disk %s was used as the sparecore disk at dumpcore time.

### Parameters

**diskname** (STRING): Name of the sparecore disk.

## coredump.spare.invalid.size

### Severity

ERROR

### Description

This message occurs when coredump subsystem encounters a spare disk with an invalid size. If an unsaved core file exists on this disk, it will be deleted because it can no longer be saved.

### Corrective Action

Run the "storage disk show -disk disk\_name" command to make sure that the disk is still connected as a spare disk. If so, replace the disk. If not, the disk is part of an aggregate or is in the process of being assigned to one, and no action is required.

### Syslog Message

sparecore disk: %s has invalid size: %llu, skipping it from all coredump subsystem operations

### Parameters

**disk** (STRING): Name of the disk.

**size** (LONGINT): Size available on the disk, reported by RAID.

## coredump.sparecore events

### coredump.sparecore.killed

### Severity

INFORMATIONAL

**Description**

Sparecore was aborted by partner.

**Corrective Action**

(None).

**Syslog Message**

Sparecore was aborted by partner.

**Parameters**

(None).

## **coredump.sparecoreinit events**

### **coredump.sparecoreinit.times**

**Severity**

NOTICE

**Description**

This message occurs when sparecore initialization is completed. Durations of constituent activities within sparecore initialization are reported.

**Corrective Action**

(None).

**Syslog Message**

sparecore\_init: update firmware %llu msec, write partial headers %llu msec, clear abort area %llu msec, write node status %lld msec.

**Parameters**

**fwUpdateTime** (LONGINT): Time taken to update firmware, in milliseconds.

**partialHeaderTime** (LONGINT): Time taken to write partial headers, in milliseconds.

**clearAbortAreaTime** (LONGINT): Time taken to clear sparecore abort area, in milliseconds.

**writeNodeStatusTime** (LONGINT): Time taken to write the node status information to the sparecore disk, in milliseconds, if any.

# cpeer events

## cpeer.access events

### cpeer.access.restricted

#### Severity

ERROR

#### Description

This message occurs when an access request from a peer cluster is rejected, either because the request is made to a restricted Vserver or to a restricted table.

#### Corrective Action

Check the Vserver peering relationships with the peer cluster. If additional Vserver access is required, use the "vserver peer create" command to initiate the peering relationship with the peer cluster.

#### Syslog Message

An access request from peer cluster %s (IP: %s) was denied. Reason: %s.

#### Parameters

**peer** (STRING): Name of the peer cluster making the request to the local cluster.

**address** (STRING): Source IP address of the incoming request.

**reason** (STRING): Reason for the access restriction.

## cpeer.addr events

### cpeer.addr.bad.dns

#### Severity

ERROR

#### Description

This message occurs when a stable address for the peer cluster cannot be resolved because of a bad name-server. This will make the peer cluster unreachable and cross-cluster operations will fail. Stable addresses are normally provided by the administrator during cluster peer creation.

#### Corrective Action

Delete the incorrect DNS server entry for the local cluster using the 'vserver services dns delete' command. Create a new DNS entry with the correct name-server using the 'vserver services dns create' command. If you need further assistance, contact NetApp technical support.

#### Syslog Message

Failed to resolve address %s of peer cluster %s, Reason:bad name-server.

#### Parameters

**addr** (STRING): Peer cluster stable IP addresses that cannot be contacted.

**peer** (STRING): Name of the peer cluster whose stable address cannot be contacted.

## **cpeer.addr.bad.host**

### **Severity**

ERROR

### **Description**

This message occurs when a stable IP address for the peer cluster is no longer valid. This will make the peer cluster unreachable and cross-cluster operations will fail. Stable addresses are normally provided by the administrator during cluster peer creation.

### **Corrective Action**

Correct the stable IP address using the 'cluster peer modify' command. If you need further assistance, contact NetApp technical support.

### **Syslog Message**

Failed to resolve address %s of peer cluster %s, Reason:invalid address.

### **Parameters**

**addr** (STRING): Peer cluster stable addresses that is not valid.  
**peer** (STRING): Name of the peer cluster whose stable address is not valid.

## **cpeer.addr.stable.down**

### **Severity**

ERROR

### **Description**

This message occurs when a stable address cannot be contacted. Stable addresses are provided by the administrator during peer creation.

### **Corrective Action**

Verify network connectivity between the local cluster and the peer cluster. Check the logs, trace buffers, and UI output for more information about the nature of the failure.

### **Syslog Message**

Stable address %s of cluster %s, in IPspace %s, cannot be contacted.

### **Parameters**

**addr** (STRING): Stable address that cannot be contacted.  
**peer** (STRING): Name of the peer cluster whose address cannot be contacted.  
**ipspace** (STRING): IPspace in which the contact attempt was made.

## **cpeer.addr.warn.host**

### **Severity**

ALERT

### **Description**

This message occurs when a stable IP address for the peer cluster is no longer valid, making the peer cluster unreachable and causing cross-cluster operations to fail.

### Corrective Action

Correct the stable IP address using the "cluster peer modify" command. Valid addresses for the remote ("peer") cluster are given under "Active IP Addresses" when you use the "cluster peer show -instance" command. Verify connectivity by using the "cluster peer ping" command. If you need further assistance, contact NetApp technical support.

### Syslog Message

Address %s is not any of the addresses that peer cluster %s considers valid for the cluster peer relationship; Details: %s

### Parameters

**addr** (STRING): Peer cluster stable address that is not valid.  
**peer** (STRING): Name of the peer cluster whose stable address is not valid.  
**complaint** (STRING): Details, if any, of the issue with the stable address.

## cpeer.authentication events

### cpeer.authentication.create

#### Severity

NOTICE

#### Description

This message occurs when authentication is successfully established with a peer cluster.

#### Corrective Action

(None).

#### Syslog Message

The peering relationship with cluster %s is now authenticated.

#### Parameters

**peer** (STRING): Peer cluster with which this cluster is now mutually authenticated.

### cpeer.authentication.disable

#### Severity

NOTICE

#### Description

This message occurs when authentication is disabled for a cluster peer relationship.

#### Corrective Action

(None).

#### Syslog Message

The peering relationship with cluster %s is now unauthenticated.

#### Parameters

**peer** (STRING): Peer cluster that is now peered to this cluster without authentication.

## cpeer.authentication.revoke

### Severity

ERROR

### Description

This message occurs when authentication is revoked for a cluster peer relationship.

### Corrective Action

Modify the relationship to match on both of the peered clusters. To have them agree on an unauthenticated relationship, use the 'cluster peer modify -authentication no-auth' command. To have them agree on an authenticated relationship, use the same passphrase with the 'cluster peer modify -authentication use-auth' command.

### Syslog Message

The peering relationship with cluster %s is now revoked and is no longer functional.

### Parameters

**peer** (STRING): Peer cluster for which the peering relationship has had its peering authentication revoked.

## cpeer.authentication.viol

### Severity

NOTICE

### Description

This message occurs when a cross-cluster access should be authenticated, but is not. The access has been denied successfully.

### Corrective Action

(None).

### Syslog Message

Access as peer cluster %s was claimed from network address %s, but has not been granted because %s.

### Parameters

**peer** (STRING): Peer cluster on behalf of which this access was made.

**remoteAddress** (STRING): Remote network address of the violating connection.

**reason** (STRING): Justification for the violation claim.

## cpeer.create events

### cpeer.create.reciprocal.fail

### Severity

ERROR

### Description

This message occurs when a peer cluster reciprocal join fails.

### **Corrective Action**

Verify network connectivity between the local cluster and the peer cluster. Check the logs, trace buffers, and UI output for more information about the nature of the failure.

### **Syslog Message**

Cluster %s attempted to create a reciprocal peer relationship with this cluster but failed.

### **Parameters**

**peer** (STRING): Peer cluster which executed the reciprocal join.

## **cpeer.create.reciprocal.ok**

### **Severity**

NOTICE

### **Description**

This message occurs when a peer cluster reciprocal join succeeds.

### **Corrective Action**

(None).

### **Syslog Message**

Cluster %s is now peered with this cluster.

### **Parameters**

**peer** (STRING): Peer cluster that executed the reciprocal join.

## **cpeer.dns events**

### **cpeer.dns.family.mismatch**

### **Severity**

ERROR

### **Description**

This message occurs when an attempt is made to resolve and insert a stable address that is a Domain Name System (DNS) style host name. If it resolves to an address family that does not match the family style specified for the cluster peer relationship, it is indicative of incorrect addresses being provided by the administrator or DNS configuration issues. Stable addresses are typically provided by the administrator during cluster peer creation and can be modified by modifying the peer relationship.

### **Corrective Action**

Correct the stable IP address by using the 'cluster peer modify' command. Ensure that LIFs of the appropriate address family exist on both the local cluster and the remote peer. Confirm that the DNS services are appropriately specified by using the 'vserver services dns show' command. Also, the DNS server in question should be configured correctly to be able to resolve the DNS style host name to an IP address belonging to the appropriate address family. If you need further assistance, contact NetApp technical support.



## Syslog Message

Host name %s of peer cluster %s resolved to %s, which does not match the peer relationship's selected address type.

## Parameters

**addr** (STRING): Peer cluster host name that is being inserted.

**peer** (STRING): Name of the peer cluster for which an address family mismatch was seen.

**resolved\_address** (STRING): Address resolved to, which belongs to the incorrect address family.

# cpeer.duplicate events

## cpeer.duplicate.remote.addr

### Severity

ERROR

### Description

This message occurs when an IP address from a peer cluster is also in use by the local cluster in the same IPspace. This may prevent the cluster peering relationship from reaching an available state.

### Corrective Action

Replace the duplicate address on the local cluster with a different address.

## Syslog Message

Cannot use address %s from peer cluster %s because it is already in use on the local cluster.";

## Parameters

**address** (STRING): Duplicate address from the peer cluster.

**peer** (STRING): Name of the peer cluster.

# cpeer.ip events

## cpeer.ip.family.mismatch

### Severity

ERROR

### Description

This message occurs when an attempt is made to insert an IP address whose address family does not match the family style specified for the cluster peer relationship. This is indicative of incorrect addresses being provided by the administrator. Stable addresses are typically provided by the administrator during cluster peer creation and can be modified by modifying the peer relationship.

### Corrective Action

Correct the stable IP address by using the 'cluster peer modify' command. Ensure that the LIFs of the appropriate address family exist on both the local cluster and the remote peer. If you need further assistance, contact NetApp technical support.

## Syslog Message

Failed to insert address %s of peer cluster %s, reason: %s

## Parameters

**addr** (STRING): Peer cluster address that is being inserted.

**peer** (STRING): Name of the peer cluster for which an address family mismatch was seen.

**reason** (STRING): Reason that the EMS message was sent.

# cpeer.localid events

## cpeer.localid.bad

### Severity

ERROR

### Description

This message occurs when there is an issue reading the local cluster ID. This indicates a serious problem with the cluster which may affect processes outside of cluster peering.

### Corrective Action

(Call support).

### Syslog Message

Unexpected local cluster ID: %s.

## Parameters

**id** (STRING): The string representation of the bad cluster ID.

# cpeer.long events

## cpeer.long.remote.update

### Severity

ERROR

### Description

This message occurs when a cross-cluster update takes longer than 15 minutes. This indicates a communication failure between the clusters.

### Corrective Action

Ensure cross-cluster network connectivity by performing network diagnostics using the "cluster peer ping" command and either the "network ping" or "network ping6" command. Correct IP address issues by using the "cluster peer modify" command. Ensure that LIFs of the appropriate address family exist on both the local cluster and the remote peer. Confirm that the DNS services are appropriately configured by using the "vserver services dns show" command. You can use "cluster peer show" to list the state of all peer clusters. This is a point-in-time command and the health of remote clusters might not have been updated yet. If you need further assistance, contact NetApp technical support.

### Syslog Message

Cross-cluster update job took longer than 15 minutes. Potential network connectivity issues.

## Parameters

(None).

# cpeer.modify events

## cpeer.modify.address.family

### Severity

NOTICE

### Description

This message occurs when the address family for a peering relationship is modified. The same configuration change should be made at the peered cluster in a timely manner. Peering services are maintained while intercluster LIFs of both address families are configured on each cluster.

### Corrective Action

This is a notification of a configuration change. No corrective action applies.

### Syslog Message

Address family change for peer cluster %s to address family %s.

### Parameters

**peer** (STRING): Name of the peer cluster affected by the address family change.

**family** (STRING): The new address family of the peering relationship.

# cpeer.out events

## cpeer.out.pp.mismatch

### Severity

ERROR

### Description

This message occurs when an outgoing cluster peer request authentication fails due to a passphrase mismatch.

### Corrective Action

For multiple occurrences of this event, verify the peering request with the admin of this local cluster. If the admin is not responsible for the attempts, the local cluster might be under attack.

### Syslog Message

This cluster did not accept a cluster peering authentication offer from cluster %s//%s due to a passphrase mismatch. Acceptance was attempted on session %s, command %s, application %s, location %s, port %s, user name %s, auth\_method %s, and PSK method %s.

### Parameters

**remotename** (STRING): Name of the cluster to which the contact was made.

**remoteuuid** (STRING): UUID of the cluster to which the contact was made.

**session** (STRING): Session ID over which the contact was made.

**command** (STRING): Command ID for which the contact was made.

**application** (STRING): Application used for the contact.

**location** (STRING): Location from which the contact was made.

**port** (STRING): Port over which the contact was made.

**username** (STRING): User name for the contact.

**authMethod** (STRING): Authentication method for the user.  
**offerStyle** (STRING): Style of authentication offer.

## cpeer.passphrase events

### cpeer.passphrase.expired

#### Severity

NOTICE

#### Description

This message occurs when the cluster peer offer for anonymous cluster with system-generated passphrase expired and deleted.

#### Corrective Action

If required, use the "cluster peer create" command with "generate-passphrase" option to generate the additional cluster peer offers.

#### Syslog Message

The cluster peer offer with system-generated passphrase (%s) expired and deleted.

#### Parameters

**passphrase\_details** (STRING): Cluster peer passphrase details.

## cpeer.ping events

### cpeer.ping.family.mismatch

#### Severity

ERROR

#### Description

This message occurs when an attempt is made to check the cluster peering relationship and no intercluster LIFs of the same address family and the correct IPspace are found on this node. Stable addresses are typically provided by the administrator during cluster peer creation and can be modified by modifying the peer relationship. The IPspace for a cluster peer relationship may be supplied when creating or modifying the relationship.

#### Corrective Action

Correct the stable IP address by using the "cluster peer modify" command. Ensure that LIFs of the appropriate address family exist on both the local cluster and the remote peer. Ensure that LIFs of the local cluster exist in the given IPspace.

#### Syslog Message

Failed to find LIFs of type %s on node %s for destination address %s in IPspace %s belonging to destination cluster %s.

#### Parameters

**peering\_family** (STRING): IP address family.  
**node** (STRING): Originating node.  
**dest\_address** (STRING): Destination IP address.

**ipspace** (STRING): IPspace for the relationship.  
**dest\_cluster** (STRING): Destination cluster.

## **cpeer.ping.no.lifs**

### **Severity**

ERROR

### **Description**

This message occurs when an attempt is made to check the cluster peering relationship and no intercluster LIFs in the correct IPspace are found on this node.

### **Corrective Action**

Create intercluster LIFs with the same address family and IPspace as the peering relationship. If the peering relationship is obsolete, remove it by using the "cluster peer delete" command.

### **Syslog Message**

Failed to find intercluster LIFs in IPspace %s on node %s.

### **Parameters**

**ipspace** (STRING): IPspace for the relationship.  
**node** (STRING): Originating node.

## **cpeer.pm events**

### **cpeer.pm.update.warn**

#### **Severity**

ERROR

#### **Description**

This message occurs when the peer manager job experiences a transient failure. The peer manager job periodically contacts one peer cluster and performs updates of cluster-wide state necessary to support cross-cluster enabled features.

#### **Corrective Action**

Verify network connectivity between the local cluster and its peer clusters. It might be necessary to manually correct the peer clusters' configurations using the 'cluster peer modify' command if peer addresses have changed recently. Verify that each cluster is up.

#### **Syslog Message**

Periodic update of peer information from %s failed. The following operations are incomplete: %s

#### **Parameters**

**peer** (STRING): Peer whose updates failed.  
**failures** (STRING): Failed peer updates.

## **cpeer.pp events**

## cpeer.pp.mismatch

### Severity

ERROR

### Description

This message occurs when an incoming cluster peer authentication request fails due to a passphrase mismatch.

### Corrective Action

For multiple occurrences of this event, verify the peering request with the admin of the remote cluster. If the admin is not responsible for the attempts, the local or remote cluster might be under attack.

### Syslog Message

Remote cluster %s//%s did not accept a cluster peering authentication offer (from address %s in IPspace %s) due to a passphrase mismatch. Local offer made on %s, valid for %s, with the %s PSK method.

### Parameters

**remotename** (STRING): Name of the cluster for which the contact was made.  
**remoteuuid** (STRING): UUID of the cluster for which the contact was made.  
**address** (STRING): Network address from which the authentication attempt was made.  
**ipspace** (STRING): IPspace from which the contact attempt was made.  
**offerTime** (STRING): Time at which the offer from the local cluster was generated.  
**offerDuration** (STRING): Duration of the offer from the local cluster.  
**offerStyle** (STRING): Style of authentication offer.

## cpeer.psk events

### cpeer.psk.cluster.revoked

### Severity

ERROR

### Description

This message occurs when the system receives an authentication request for a remote cluster whose peer authentication has been revoked by the local cluster administrator. It might indicate that this system is under attack.

### Corrective Action

ONTAP® software automatically closes the connection from which this incorrect authentication information was received. If this message is repeatedly generated from the same internet address, it might indicate that this system is under attack.

### Syslog Message

ONTAP received a TLS request to authenticate with the pre-shared key of remote cluster %s, but the relationship with the remote cluster has been revoked by the administrator of this cluster. The request was sent from address %s in IPspace %s.

### Parameters

**clustername** (STRING): Name of the remote cluster being probed.  
**sourceIP** (STRING): IP address that is the source of the probe.

**sourceIPspace** (STRING): IPspace in which the probe arrived.

## **cpeer.psk.cluster.unauth**

### **Severity**

ERROR

### **Description**

This message occurs when the system receives an authentication request for a remote cluster whose peer relationship is unauthenticated. It might indicate that this system is under attack. The request was sent from address %s in IPspace %s.

### **Corrective Action**

ONTAP® software automatically closes the connection from which this incorrect authentication information was received. If this message is repeatedly generated from the same internet address, it might indicate that this system is under attack.

### **Syslog Message**

ONTAP received a TLS request to authenticate with a cluster's pre-shared key, but there is no authenticated relationship with this remote cluster. Instead, the relationship is unauthenticated and has no keying information associated with it.

### **Parameters**

**clustername** (STRING): Name of the remote cluster being probed.

**sourceIP** (STRING): IP address that is the source of the probe.

**sourceIPspace** (STRING): IPspace in which the probe arrived.

## **cpeer.psk.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the system is unable to find a suitable pre-shared key (PSK) cipher suite required for cluster peering. This is likely because the PSK cipher suites are disabled in the ONTAP security configuration. Cluster peering connections and associated applications may fail as a result.

### **Corrective Action**

View the enabled cipher suites with the "security config show" command. Enable the PSK cipher suites with the "security config modify" command.

### **Syslog Message**

The system was unable to find a suitable pre-shared key (PSK) cipher suite required for cluster peering. This is likely because the PSK cipher suites are disabled in the ONTAP security configuration. Cluster peering connections and associated applications may fail as a result.

### **Parameters**

(None).

## **cpeer.psk.unknown.cluster**

## Severity

ERROR

## Description

This message occurs when the system receives an authentication request from an unknown cluster. It might indicate that this system is under attack.

## Corrective Action

ONTAP® software automatically closes the connection from which this incorrect authentication information was received. If this message is repeatedly generated from the same internet address, it might indicate that this system is under attack.

## Syslog Message

ONTAP received a TLS request to authenticate with a remote cluster's pre-shared key, but the remote cluster UUID (%s) is unrecognized. The probe arrived from address %s in IPspace %s.

## Parameters

**clusterUUID** (STRING): Cluster UUID being probed.

**sourceIP** (STRING): IP address that is the source of the probe.

**sourceIPspace** (STRING): IPspace in which the probe arrived.

# cpeer.receiver events

## cpeer.receiver.noindex

### Severity

NOTICE

### Description

This message occurs when an incoming cluster peer request using the system-generated passphrase. However, the cluster peer request failed because there is no matching outstanding cluster peer offers for the request in the system.

### Corrective Action

Verify the peer request with peer cluster admin. If required, use the "cluster peer create" command with "generate-passphrase" option to generate the additional cluster peer offers.

### Syslog Message

The cluster peer request from the remote cluster (%s) using the system-generated passphrase (Index: %d) has failed because cluster peer offer with system-generated passphrase is not available for this index.

### Parameters

**remote\_cluster\_details** (STRING): Remote cluster name.

**ppindex** (INT): Cluster peer passphrase index.

# cpeer.recv events

## cpeer.recv.pp.mismatch



**Severity**

NOTICE

**Description**

This message occurs when there is an incoming cluster peer request using the system-generated passphrase. However, the cluster peer request authentication failed due to passphrase mismatch and so the cluster peer offer is deleted.

**Corrective Action**

Verify the peer request with peer cluster admin. If required, use the "cluster peer create" command with "generate-passphrase" option to generate the additional cluster peer offers.

**Syslog Message**

The cluster peer offer with system-generated passphrase was attempted for cluster peering by the remote cluster (%s:%s) with address %s but cluster peer request failed due to passphrase mismatch. So, the offer is deleted.

**Parameters**

**remotename** (STRING): Remote cluster name.  
**remoteuuid** (STRING): Remote cluster UUID.  
**address** (STRING): Remote cluster IP address.

## cpeer.unavailable events

### cpeer.unavailable

**Severity**

ALERT

**Description**

This message occurs when a peer cluster becomes unavailable.

**Corrective Action**

Verify network connectivity between the local cluster and the peer cluster. It might be necessary to manually correct the peer configuration using the 'cluster peer modify' command if peer addresses have changed recently. Verify that the peer cluster is up.

**Syslog Message**

Peer cluster %s is no longer available.

**Parameters**

**peer** (STRING): Peer cluster that became unavailable.

## cpeer.xcm events

### cpeer.xcm.addr.disc.warn

**Severity**

ERROR

## Description

This message occurs when address discovery fails for a peer cluster.

## Corrective Action

Verify network connectivity between the local cluster and the peer cluster. It might be necessary to manually correct the peer configuration using the 'cluster peer modify' command if peer addresses have changed recently. Verify that the peer cluster is up.

## Syslog Message

Address discovery failed for peer cluster %s. Reason: %s.

## Parameters

**peer** (STRING): Peer cluster whose address discovery failed.

**reason** (STRING): Reason for address discovery failure.

## cpeer.xcm.update.warn

### Severity

ERROR

### Description

This message occurs when the cross-cluster manager job experiences a transient failure. The cross-cluster manager job periodically contacts each peer cluster, requests configuration information, and maintains the node-local network configuration for the remote peer cluster.

### Corrective Action

Verify network connectivity between the local cluster and its peer clusters. It might be necessary to manually correct the peer clusters' configurations using the 'cluster peer modify' command if peer addresses have changed recently. Verify that each cluster is up.

### Syslog Message

Periodic update of peer network information failed. The following operations are incomplete: %s.

### Parameters

**failures** (STRING): Failed peer updates.

# cr events

## cr.data events

### CR.Data.File.Inaccessible

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

NOTICE

#### Description

This message occurs when the namespace scanner skips an inode for which it cannot access the data file. This might be due to the data volume being offline.

#### Corrective Action

Bring the data volume back online using the "volume online" command. If the problem persists, retry the scan operation.

#### Syslog Message

The redirector file [%s] on namespace volume [%s] will be skipped by "Volume Check" because the corresponding data file on data volume [%s] in Vserver [%s] is not accessible.

#### Parameters

**redirPath** (STRING): Path of the redirector file that will be skipped.

**nsVol** (STRING): Name of the constituent namespace volume that hosts the redirector file.

**dataVol** (STRING): Name of the constituent data volume that hosts the data file.

**vserver** (STRING): Vserver hosting the constituent volumes.

## cr.del events

### CR.Del.CrptStreamData.Fail

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

NOTICE

#### Description

This message occurs when an attempt to delete corrupt stream data fails.

#### Corrective Action

(None).

#### Syslog Message

Deletion failed for corrupt stream data [%s] on volume [%s] in Vserver [%s].

## Parameters

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the stream resides.

**vserver** (STRING): Vserver hosting the volume.

## CR.Del.CrptStreamRedir.Fail

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when an attempt to delete corrupt stream redirector failed.

### Corrective Action

(None).

### Syslog Message

Deletion failed for corrupt stream redirector [%s] on volume [%s] in Vserver [%s].

## Parameters

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the stream resides.

**vserver** (STRING): Vserver hosting the volume.

## CR.Del.DangStreamData.Fail

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when an attempt to delete dangling stream data fails. This stream was created when there was a failure in cleaning up stream data on the data constituent, which left dangling data that is unnecessarily consuming resources.

### Corrective Action

(None).

### Syslog Message

Deletion failed for dangling stream data [%s] on volume [%s] in Vserver [%s].

## Parameters

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the stream data resides.

**vserver** (STRING): Vserver hosting the volume.

## CR.Del.DangStreamRedir.Fail

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when an attempt to delete dangling stream redirector failed.

### Corrective Action

(None).

### Syslog Message

Deletion failed for dangling stream redirector [%s] on volume [%s] in Vserver [%s].

### Parameters

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the redirector resides.

**vserver** (STRING): Vserver hosting the volume.

## cr.fix events

### CR.Fix.Corrupt.Redir.Failed

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

NOTICE

#### Description

This message occurs when corrective action for corrupt redirector file failed.

#### Corrective Action

(None).

#### Syslog Message

Corrective action failed for corrupt redirector of data file [%s] on volume [%s] in Vserver [%s].

#### Parameters

**path** (STRING): File path.

**volume** (STRING): Name of the constituent volume where the data file resides.

**vserver** (STRING): Vserver hosting the volume.

### CR.Fix.Crpt.Data.Dir.Failed

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

INFORMATIONAL

**Description**

This message occurs when the corrective action for a corrupt data directory fails. These entries were created when they were moved to the [lost+found] folder by wafllron, on the data constituents.

**Corrective Action**

(None).

**Syslog Message**

Corrective action failed for corrupt data directory [%s] on volume [%s].

**Parameters**

**path** (STRING): Directory path.

**volume** (STRING): Name of the constituent volume where the directory resides.

**vserver** (STRING): Vserver hosting the volume.

**CR.Fix.Crpt.Data.File.Failed****Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when the corrective action for a corrupt data file fails.

**Corrective Action**

(None).

**Syslog Message**

Corrective action failed for corrupt data file [%s] on volume [%s] in Vserver [%s].

**Parameters**

**path** (STRING): File path.

**volume** (STRING): Name of the constituent volume where the file resides.

**vserver** (STRING): Vserver hosting the volume.

**CR.Fix.CrptStreamRedir.Fail****Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when the fix for corrupt stream redirector fails.

**Corrective Action**

(None).

**Syslog Message**

Corrective action failed for corrupt stream redirector of stream data [%s] on volume [%s] in Vserver [%s].

**Parameters**

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the stream redirector resides.

**vserver** (STRING): Vserver hosting the volume.

**CR.Fix.Dang.Data.File.Failed****Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when the corrective action for a dangling data file fails. This file was created when there was a failure in cleaning up a file on the data constituent, which left a dangling file that is unnecessarily consuming resources.

**Corrective Action**

(None).

**Syslog Message**

Corrective action failed for dangling data file [%s] on volume [%s] in Vserver [%s].

**Parameters**

**path** (STRING): File path.

**volume** (STRING): Name of the constituent volume where the file resides.

**vserver** (STRING): Vserver hosting the volume.

**CR.Fix.Nlinks.Failed****Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when the corrective action for an Nlinks (number of hardlinks on the file) mismatch fails.

**Corrective Action**

(None).

## Syslog Message

Corrective action failed for Nlinks mismatch on file [%s] on volume [%s] in Vserver [%s].

## Parameters

**path** (STRING): File path.

**volume** (STRING): Name of the constituent volume where the file resides.

**vserver** (STRING): Vserver hosting the volume.

# cr.rdb events

## CR.RDB.Counters.Not.Updated

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when the RDB (replicated database) cannot be updated with the inode counters for a particular scan.

### Corrective Action

(None).

## Syslog Message

The inode counters [percentage scan complete: %d], [estimated time remaining: %s], [number of issues detected: %d] and [number of issues fixed: %d] could not be updated in the RDB due to [error: %s] for [volume: %s] in [Vserver: %s].

## Parameters

**percentComplete** (INT): Percentage of the inodes scanned.

**estimatedTime** (STRING): Estimated time remaining for the scan to finish.

**issuesDetected** (INT): Number of issues detected by the scan so far.

**issuesFixed** (INT): Number of issues fixed by the scan so far.

**reason** (STRING): Reason for the failed update.

**volume** (STRING): Name of the constituent volume for which the RDB could not be updated.

**vserver** (STRING): Vserver hosting the volume.

## CR.RDB.State.Not.Updated

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when the RDB (replicated database) cannot be updated with the final state of the scan, that is, success/failure.



## Corrective Action

Run the 'volume check stop' command on the respective volume.

## Syslog Message

The final state [%s] of the 'volume check' scan could not be updated in the RDB due to [error: %s] for volume [%s] in Vserver [%s].

## Parameters

**state** (STRING): Final state of the scan that has to be updated.

**reason** (STRING): Reason for the failed update.

**volume** (STRING): Name of the constituent volume for which the RDB could not be updated.

**vserver** (STRING): Vserver hosting the volume.

# cr.redir events

## CR.Redir.File.Inaccessible

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when the data scanner skips an inode for which it cannot access the redirector file. This might be due to the namespace volume being offline.

### Corrective Action

Bring the namespace volume back online using the "volume online" command. If the problem persists, retry the scan operation.

### Syslog Message

The data file [%s] on data volume [%s] will be skipped by "Volume Check" because the corresponding redirector file on namespace volume [%s] on Vserver [%s] is not accessible.

### Parameters

**dataFilePath** (STRING): Path of the data file that will get skipped.

**dataVol** (STRING): Name of the constituent data volume that hosts the data file.

**nsVol** (STRING): Name of the constituent namespace volume that hosts the redirector file.

**vserver** (STRING): Vserver hosting the constituent volumes.

# cr.snapshot events

## CR.Snapshot.Not.Deleted

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when a temporary Snapshot(tm) copy created by the 'volume check' command could not be deleted.

**Corrective Action**

Delete the Snapshot copy by using the following command from "diag" mode: volume snapshot delete -vserver <vserver> -volume <volume> -snapshot <snapshot-name> -ignore-owners

**Syslog Message**

The temporary Snapshot copy [%s] created by the 'volume check' command for volume [%s] in Vserver [%s] could not be deleted [error: %s].

**Parameters**

**snapshot** (STRING): Name of the Snapshot copy.

**volume** (STRING): Name of the constituent volume for which the Snapshot copy could not be deleted.

**vserver** (STRING): Vserver hosting the volume.

**reason** (STRING): Reason for the failure.

## cr.sync events

### CR.Sync.ACL.Fail

**Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

NOTICE

**Description**

This message occurs when an attempt to synchronize the ACL from data to redirector fails.

**Corrective Action**

(None).

**Syslog Message**

ACL synchronizing failed for file [%s] on volume [%s] in Vserver [%s].

**Parameters**

**path** (STRING): Stream path.

**volume** (STRING): Name of the constituent volume where the file resides.

**vserver** (STRING): Vserver hosting the volume.

# crypto events

## crypto.debug events

### crypto.debug

**Severity**

INFORMATIONAL

**Description**

The Crypto module debug event.

**Corrective Action**

(None).

**Syslog Message**

%s.

**Parameters**

**debug\_string** (STRING): Description of the event being logged.

## crypto.export events

### crypto.export.failed

**Severity**

ALERT

**Description**

This message occurs when the export of a cryptographic key fails.

**Corrective Action**

This message might be seen when attempting to create an encrypted volume. Contact NetApp technical support if further attempts to create the encrypted volume fail.

**Syslog Message**

ERROR: Export of key with key ID %s failed. Additional information: %s.

**Parameters**

**key\_id** (STRING): Key ID for the cryptographic key.

**msg** (STRING): Associated failure message.

## crypto.import events

### crypto.import.failed

**Severity**

ALERT

## Description

This message occurs when the import of a cryptographic key fails.

## Corrective Action

This message might be seen when attempting to bring an encrypted volume online. Contact NetApp technical support if manual attempts to bring the encrypted volume online fail.

## Syslog Message

ERROR: Import of key with key ID %s failed. Additional information: %s.

## Parameters

**key\_id** (STRING): Key ID for the cryptographic key.

**msg** (STRING): Associated failure message.

# crypto.key events

## crypto.key.deleted

### Severity

NOTICE

### Description

This message occurs when a cryptographic key is deleted.

### Corrective Action

(None).

### Syslog Message

Deleted key with key ID %s. Reason %s.

### Parameters

**key\_id** (STRING): Key ID for the cryptographic key.

**reason** (STRING): Reason for deletion.

## crypto.key.stored

### Severity

NOTICE

### Description

This message occurs when a cryptographic key is stored.

### Corrective Action

(None).

### Syslog Message

Stored key with key ID %s. Key digest: %s.

### Parameters

**key\_id** (STRING): Key ID for the cryptographic key.

**key\_digest** (STRING): Key digest for the cryptographic key.

## crypto.okmrecovery events

### crypto.okmrecovery.failed

#### Severity

ALERT

#### Description

This message occurs when the Onboard Key Manager (OKM) recovery command in the boot menu fails to recover the OKM hierarchy. Without the the OKM hierarchy, volumes encrypted using OKM will not come online, and you will not be able to use OKM to encrypt volumes.

#### Corrective Action

This message is usually seen when an incorrect cluster passphrase was entered during a previous boot menu recovery command. Often, using the "security key-manager onboard sync" command will resolve the situation. If that does not fix the problem, the OKM hierarchy can be recovered during the boot process. To recover the OKM hierarchy during the boot process, enter option 10 ("Set Onboard Key Manager recovery secrets.") at the boot menu prompt. You will need to have the cluster passphrase and a copy of the output from the "security key-manager backup show" command. If the boot menu recovery process fails again, contact NetApp technical support.

#### Syslog Message

Import of the Onboard Key Manager (OKM) hierarchy has failed: %s. Additional information: %s.

#### Parameters

**failure\_msg** (STRING): Message describing which operation failed.

**additional\_msg** (STRING): Additional information regarding the failure.

## crypto.ssal events

### crypto.ssal.failed

#### Severity

ALERT

#### Description

This message occurs when a Secure Storage Access Layer(SSAL) operation fails. This might cause import of the onboard key hierarchy to fail and might result in a temporary loss of access to any data secured using the SSAL.

#### Corrective Action

(Call support).

#### Syslog Message

SSAL operation failed: %s. %s

#### Parameters

**failure\_msg** (STRING): Message describing which operation failed.

**additional\_msg** (STRING): Additional information regarding the failure.

## **crypto.ssal.tpm.clear**

### **Severity**

NOTICE

### **Description**

This message occurs when a Trusted Platform Module(TPM) clear operation is performed successfully. A TPM clear operation can happen while reverting to a previous release or when the "security tpm clear" command was executed. A TPM clear operation clears out the TPM and resets it with a new seed for the TPM's storage hierarchy.

### **Corrective Action**

(None).

### **Syslog Message**

A TPM clear operation has been completed successfully.

### **Parameters**

(None).

## **crypto.ssal.tpm.reset**

### **Severity**

NOTICE

### **Description**

This message occurs when a Trusted Platform Module(TPM) reset operation is performed successfully. A TPM reset operation can happen when the node boots up for the first time or when the "security tpm reset" command was executed. A TPM reset operation clears out the TPM and resets it with a new seed for the TPM's storage hierarchy. The TPM then generates a new primary key under the same hierarchy and makes it persistent.

### **Corrective Action**

(None).

### **Syslog Message**

A TPM reset operation has been completed successfully.

### **Parameters**

(None).

# cs events

## cs.create events

### cs.create.inprogress

#### Severity

ALERT

#### Description

This message occurs when there is an attempt to release SnapMirror resources for a volume as a result of volume delete while a Cloud Snapshot create operation is in progress.

#### Corrective Action

Ensure no Cloud Snapshot create operation is in-progress when the volume is deleted. Either abort the create operation or wait for the creation to complete before deleting the volume.

#### Syslog Message

Failed to delete the volume "%s" as one or more Cloud Snapshot operations are in progress. Number of Cloud Snapshot operations in progress: "%d".

#### Parameters

**volumeUUID** (STRING): UUID of the volume.

**in\_transfer\_snap\_count** (INT): Number of Cloud Snapshot create in-progress.

## cs.endpoint events

### cs.endpoint.has.no.snapshots

#### Severity

ALERT

#### Description

This message occurs when volume delete operations attempts to clean-up snapmirror resources and finds that there are no cloud snapshots associated with the volume. The cloud endpoint is automatically deleted on deletion of last cloud snapshot. So it's unusual to find a cloud snapshot endpoint with no snapshots.

#### Corrective Action

(None).

#### Syslog Message

The volume "%s" has no cloud snapshots associated with it.

#### Parameters

**volumeUUID** (STRING): UUID of the volume.

# csm events

## csm.badauthaccess events

### csm.badAuthAccess

#### Severity

NOTICE

#### Description

This message occurs when a blade fails an attempt to authenticate as a given peer cluster, and Cluster Session Manager (CSM), which establishes and manages SpinNP sessions between blades in the cluster, denies access of the unauthenticated blade. This could indicate that an interloper is attempting to gain access to the cluster.

#### Corrective Action

(None).

#### Syslog Message

Access as peer cluster %s was claimed from network address %s, but has not been granted because %s.

#### Parameters

**clusterUuid** (STRING): UUID of the cluster to which the blade claimed to belong.

**addrstr** (STRING): Remote network address of the blade.

**msg** (STRING): Reason that access could not be authenticated.

## csm.badconnection events

### csm.badConnection

#### Severity

ALERT

#### Description

This message occurs when ONTAP® software receives a Cluster Session Manager (CSM) connection with unrecognizable content. This might indicate that the system is under attack. As a precaution, the connection used to receive that data has been closed.

#### Corrective Action

Carefully inspect the networking configuration, event messages, and logs of both the local and remote nodes. If any evidence exists of an intrusion attempt, configure site networking equipment to avoid forwarding data to the local address and port, and if necessary, from forwarding data from the remote address and port. Check external site security tools for unauthorized access attempts. Audit site security logs.

#### Syslog Message

ONTAP received a CSM connection with unrecognizable content at local address %s local port %d, from remote address %s remote port %d, via IPspace %d.



## Parameters

**localAddr** (STRING): Full internet address of the local end of the connection.  
**localPort** (INT): Local port number.  
**remoteAddr** (STRING): Full internet address of the remote end of the connection.  
**remotePort** (INT): Remote port number.  
**IPspace** (INT): Identifier of the IPspace within which the remote address was reached.

# csm.badplainaccess events

## csm.badPlainAccess

### Severity

NOTICE

### Description

This message occurs when a blade attempts to access the local cluster without encrypting its communication channel, and Cluster Session Manager (CSM), which establishes and manages SpinNP sessions between blades in the cluster, denies access of the unencrypted communication channel. This could indicate that an interloper is attempting to gain access to the cluster.

### Corrective Action

(None).

### Syslog Message

Access as peer cluster %s was claimed from network address %s, but has not been granted because it was not using TLS.

### Parameters

**clusterUuid** (STRING): UUID of the cluster to which the blade claimed to belong.  
**addrstr** (STRING): Remote network address of the blade.

# csm.ctfallbackactiveopen events

## csm.ctFallbackActiveOpen

### Severity

NOTICE

### Description

This message occurs when the Cluster Session Manager (CSM) uses CT (TCP transport) as the default when creating LIF pairs because connections running CSM over remote direct memory access (RDMA) repeatedly fail during successive retry attempts.

### Corrective Action

(None).

### Syslog Message

Cluster Session Manager (CSM) could not successfully create the RDMA connections for session "%s" even after several retry attempts. CSM will use TCP connections as defaults.

## Parameters

**uniquifier** (STRING): Unique identifier for this session.

# csm.ctfallbackswitchless events

## csm.ctFallbackSwitchless

### Severity

NOTICE

### Description

This message occurs when the Cluster Session Manager (CSM) defaults to use CT (TCP transport) when creating LIF pairs because remote device ID information is not available on switchless clusters running CSM over remote direct memory access (RDMA).

### Corrective Action

(None).

### Syslog Message

Cluster Session Manager (CSM) could not determine the cluster ports that are directly connected between the switchless cluster nodes to create the RDMA connections for session "%s". CSM will use TCP connections as defaults.

## Parameters

**uniquifier** (STRING): Unique identifier for this session.

# csm.homeportinfounexpected events

## csm.homePortInfoUnexpected

### Severity

NOTICE

### Description

This message occurs when the Cluster Session Manager (CSM) establishes a connection between nodes over the cluster network interface of a switchless cluster and detects that the logical interface's (LIF's) home port and the corresponding remote device information are not the expected values.

### Corrective Action

This message can occasionally happen during normal operations on a system startup or a LIF migration event. If the message repeats every few minutes, it might indicate a problem in the cluster interface configuration or the physical cabling of the switchless cluster.

### Syslog Message

CSM reports unexpected information for cluster LIF ID "%d" (home port: "%s", remote device ID: "%s").

## Parameters

**vifId** (INT): ID of the local cluster LIF over which the connection gets established.

**homePort** (STRING): Home port hosting the LIF.

**remoteDeviceId** (STRING): Remote device associated with the home port.

# csm.mismatchremotedevice events

## csm.mismatchRemoteDevice

### Severity

ERROR

### Description

The message occurs when the Cluster Session Manager (CSM) establishes a connection between nodes over the cluster network interface, but the node's remote device IDs do not match.

### Corrective Action

Ensure that the Cisco Discover Protocol (CDP) is running on the nodes and switches. In addition, ensure that the cluster ports are up and the cluster LIFs are configured and hosted according to the suggested cluster configuration.

### Syslog Message

CSM connection between source LIF %d and destination address %s might not be optimal for session %s. The source is currently connected to %s remote device and the destination is currently connected to %s remote device.

### Parameters

**sourceLIF** (INT): Source logical interface.

**destAddr** (STRING): Destination IP address.

**uniquifier** (STRING): Unique identifier for this session.

**sourceRemoteDevice** (STRING): Remote device associated with the source logical interface.

**destRemoteDevice** (STRING): Remote device associated with the destination IP address.

# csm.nomemory events

## csm.noMemory

### Severity

ERROR

### Description

This message occurs when the memory allocator temporarily runs out of memory that can be allocated or when the entire node has an insufficient amount of free system memory. If the issue persists, the cluster session manager (CSM) might not be able to create connections to remote cluster nodes to serve data traffic.

### Corrective Action

If the condition persists, contact NetApp technical support.

### Syslog Message

CSM was unable to process events due to insufficient free memory.

### Parameters

**filename** (STRING): Name of the file where the event occurred.

**lineno** (INT): Line number where the event occurred.

## **csm.rdmarecoveractiveopen events**

### **csm.rdmaRecoverActiveOpen**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the Cluster Session Manager (CSM) attempts to restore the use of the preferred remote direct memory access (RDMA) transport for the eligible sessions because the condition which triggered the CSM to use CT (TCP transport) as default might have resolved.

#### **Corrective Action**

(None).

#### **Syslog Message**

Cluster Session Manager (CSM) is attempting to restore the use of the preferred RDMA connections for eligible sessions, which had previously defaulted to use TCP connections due to repeated active open failures.

#### **Parameters**

(None).

## **csm.rdmarecoverswitchless events**

### **csm.rdmaRecoverSwitchless**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the Cluster Session Manager (CSM) attempts to restore the use of the preferred remote direct memory access (RDMA) transport for the eligible sessions because the condition which triggered the CSM to use CT (TCP transport) as default might have been resolved in the switchless cluster.

#### **Corrective Action**

(None).

#### **Syslog Message**

Cluster Session Manager (CSM) is attempting to restore the use of the preferred RDMA connections for eligible sessions, which had previously defaulted to use TCP connections in the switchless cluster.

#### **Parameters**

(None).

## **csm.unmarshallerror events**

## **csm.unmarshallError**

### **Severity**

ERROR

### **Description**

The message occurs when the Cluster Session Manager (CSM) fails to unmarshal a packet.

### **Corrective Action**

Unmarshal errors reported in CSM indicate corrupted packet(s) were delivered to CSM. Often this can be attributed to network layer problems. TCP checksum errors should be checked for if unmarshal errors continue to increment.

### **Syslog Message**

Cluster Session Manager (CSM) failed to unmarshall packet for session %s.

### **Parameters**

**sessionId** (STRING): Full session ID of this session.

# ctl events

## ctl.ctlbe events

### ctl.ctlbe.IOInternalError

**Severity**

ERROR

**Description**

IO failure due to internal error

**Corrective Action**

The command failure associated with the device specified in this event is caused by some internal error.

**Syslog Message**

Device Serial Number %s - Internal error reason: %s - SCSI error : %s

**Parameters**

**devSerial** (STRING): Device Serial Number

**reason** (STRING): Internal error

**scsiError** (STRING): SCSI error code

### ctl.ctlbe.lunCreationFailed

**Severity**

ERROR

**Description**

LUN creation failed in the CAM target layer.

**Corrective Action**

The creation of a CAM target end point for the LUN failed. This will cause the disk or partition to not be exported to the DR cluster, and therefore cannot be used for storage mirroring.

**Syslog Message**

LUN %d - endpoint creation failed, reason: %s

**Parameters**

**lunId** (INT): The LUN ID

**errorString** (STRING): Description of the error

### ctl.ctlbe.lunError

**Severity**

ERROR

**Description**

The LUN has been deleted or is not configured properly

## Corrective Action

The command failure associated with the device specified in this event is caused by the device being not configured or in error state. The device may be in the process of being deleted or the path to this device may not be accessible. Check the state of the device.

## Syslog Message

Device Serial Number %s - Device is not configured or in error state - SCSI error : %s - LUN flags : 0x%x

## Parameters

**devSerial** (STRING): Device Serial Number  
**scsiError** (STRING): SCSI error code  
**lunFlags** (INT): Internal CTL Backend LUN flags

## ctl.ctlbe.lunNotConfigured

### Severity

ERROR

### Description

The LUN is not configured

### Corrective Action

The command failure associated with the device specified in this event is caused by the device not being configured. Check the state of the device.

## Syslog Message

Device Serial Number %s - Device is not configured - SCSI error : %s

## Parameters

**devSerial** (STRING): Device Serial Number  
**scsiError** (STRING): SCSI error code

## ctl.session events

### ctl.session.stateChanged

### Severity

NOTICE

### Description

This message occurs when iSCSI CAM target layer session status is changed.

### Corrective Action

If the identified iSCSI CAM target layer session is experiencing unexpected events (disconnecting, reconnecting, or connected), examine the link for excessive errors or latency.

## Syslog Message

iSCSI CAM target layer's session state is changed to %s for the initiator %s (address: %s). %s

## Parameters

**ctl\_session\_state** (STRING): New iSCSI CAM target layer session state.

**ctl\_initiator** (STRING): iSCSI CAM target layer initiator's name.

**ctl\_initiator\_address** (STRING): iSCSI CAM target layer initiator's address.

**ctl\_session\_state\_change\_reason** (STRING): Reason for the iSCSI CAM target layer session state change.

**ctl\_timeout\_value** (INT): iSCSI CAM target layer timeout value.



# ctran events

## ctran.disabled events

### ctran.disabled

#### Severity

INFORMATIONAL

#### Description

This message occurs when Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN) subsystem is disabled.

#### Corrective Action

This happens only when the local node went out of CLAM quorum. Take the corrective action specified in EMS\_clam\_node\_ooq event.

#### Syslog Message

CTRAN subsystem is now disabled.

#### Parameters

(None).

## ctran.enabled events

### ctran.enabled

#### Severity

INFORMATIONAL

#### Description

This message occurs when Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN) subsystem is enabled.

#### Corrective Action

No action is needed.

#### Syslog Message

CTRAN subsystem is now enabled.

#### Parameters

(None).

## ctran.gbuf events

### ctran.gbuf.alloc.unsuccessful

## Severity

NOTICE

## Description

This message occurs when one or more attempts to allocate a message buffer to convey messages between nodes for the Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN) has been unsuccessful.

## Corrective Action

If this message occurs more than 10 times in a hour KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using, "autosupport invoke -node \* -message "ctran.gbuf.alloc.unsuccessful" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.gbuf.alloc.unsuccessful" -subsystem kcs" and contact NetApp technical support support.

## Syslog Message

Attempt to allocate a message buffer for CTRAN has been unsuccessful. Total number of unsuccessful attempts is %u.

## Parameters

**attempts\_failed** (LONGINT): Total number of message buffer allocation attempts that have been unsuccessful.

# ctran.gc events

## ctran.gc.proposal.stuck

## Severity

EMERGENCY

## Description

This message occurs when the Cluster Transaction Manager (CTRAN) Group Coordinator (GC) of a group runs a prop-done probe for a proposal longer than anticipated.

## Corrective Action

KCS CTRAN logs need to be collected for further analysis. Trigger an ALL AutoSupport® message by using the "autosupport invoke -node \* -message "ctran.gc.proposal.stuck" -type all" command. Also trigger a diagnostic AutoSupport by using the "autosupport invoke-diagnostic -node \* -message "ctran.gc.proposal.stuck" -subsystem kcs" command. Then contact NetApp technical support support.

## Syslog Message

CTRAN group %s has attempted a prop\_done probe %d times against gm\_id %ld on node %s with UUID %s for txn %ld with opcode %d.

## Parameters

**group\_name** (STRING): Name of the group being polled.  
**attempts** (LONGINT): Number of attempts.  
**gm\_id** (LONGINT): Group member (GM) ID.  
**node\_name** (STRING): Name of the node the GM resides on.  
**node\_uuid** (STRING): UUID of the node the GM resides on.  
**trans\_id** (LONGINT): ID of the stuck transaction.  
**opcode** (INT): Operation type of the stuck transaction.

## ctran.gc.reselection.failed

### Severity

ALERT

### Description

This message occurs when some Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN) groups lose their Group co-ordinator (GC) and a subsequent GC reselection process fails to complete.

### Corrective Action

Reboot the node, and if the node is waiting for giveback then perform a giveback.

### Syslog Message

CTRAN GC Reselection failed on local node (name=%s, id=%llu).

### Parameters

**localhostname** (STRING): The host name of the local node.

**localnodeid** (LONGINT): The node ID of the local node.

## ctran.group events

### ctran.group.reset.failed

### Severity

EMERGENCY

### Description

This message occurs when an attempt to reset a group by Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN) fails.

### Corrective Action

KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using "autosupport invoke -node \* -message "ctran.group.reset.failed" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.group.reset.failed" -subsystem kcs" and contact NetApp technical support support.

### Syslog Message

Group reset for %s has failed on node %s.

### Parameters

**localhostname** (STRING): The host name of the local node.

**group\_name** (STRING): The group on which the group reset operation is attempted.

## ctran.jpc events

### ctran.jpc.assigned

### Severity

INFORMATIONAL

## Description

This message occurs when a Join Proposal Coordinator (JPC) node is assigned by Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN).

## Corrective Action

(None).

## Syslog Message

Cluster node (name=%s, ID=%llu) is the Join Proposal Coordinator (JPC) node.

## Parameters

**nodename** (STRING): Host name of the JPC node.

**node\_id** (LONGINT): Node id of the JPC node.

## ctran.jpc.multiple.nodes

### Severity

EMERGENCY

### Description

This message occurs when multiple Join Proposal Coordinator (JPC) nodes for the Cluster Transaction Manager (CTRAN) are detected.

### Corrective Action

KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using, "autosupport invoke -node \* -message "ctran.jpc.multiple.nodes" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.jpc.multiple.nodes" -subsystem kcs" and contact NetApp technical support support.

### Syslog Message

Cluster node %u detected multiple JPC nodes (node %u reports JPC %u, node %u reports JPC %u)

### Parameters

**local\_node\_id** (LONGINT): ID of the local node.

**src\_node\_id\_1** (LONGINT): ID of the first node reporting a JPC.

**jpc\_node\_id\_1** (LONGINT): ID of the JPC node reported by src\_node\_id\_1.

**src\_node\_id\_2** (LONGINT): ID of the second node reporting a JPC.

**jpc\_node\_id\_2** (LONGINT): ID of the JPC node reported by src\_node\_id\_2.

## ctran.jpc.not.found

### Severity

EMERGENCY

### Description

This message occurs when this node cannot find a a Join Proposal Coordinator (JPC) within the cluster for Kernel Cluster Services (KCS) Cluster Transaction Manager (CTRAN). Until this error is resolved the reporting node will not be able to participate in distributed operations.

### Corrective Action

Perform a planned takeover/giveback sequence to reset the reporting node non-disruptively by using the "takeover -ofnode (reporting node)" and "giveback -ofnode (reporting node)" commands. If the error

message is still seen KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using, "autosupport invoke -node \* -message "ctran.jpc.not.found" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.jpc.not.found" -subsystem kcs" and contact NetApp technical support support.

### Syslog Message

This node was unable to locate a Join Proposal Coordinator (JPC) node.

### Parameters

(None).

## ctran.jpc.valid.failed

### Severity

EMERGENCY

### Description

This message occurs when the Join Proposal Coordinator (JPC) validation fails due to an inconsistent state of some of the Cluster Transaction Manager (CTRAN) group members following a JPC reselection event. In some cases, this can result in loss of SAN data access, CLI command failures and/or giveback failures.

### Corrective Action

Reset the CTRAN subsystem during a maintenance window by placing the node specified ineligible by using "cluster modify -node (node reporting error) -eligibility false". Once complete return the node to service by entering "cluster modify -node (node reporting error) -eligibility true". Note, making a node ineligible impacts all protocols on that node. During the time that a node is ineligible, it will not serve any data. If a non-disruptive resolution is desired, takeover the reporting node by using "takeover -of-node (reporting-node)", and then returning it to service by "giveback -of-node (reporting-node)" when it is ready. If the error is still seen KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using "autosupport invoke -node \* -message "ctran.jpc.valid.failed" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.jpc.valid.failed" -subsystem kcs" and contact NetApp technical support support.

### Syslog Message

CTRAN JPC validation failed on local node (name=%s, id=%llu, num\_groups=%llu).

### Parameters

**localhostname** (STRING): Host name of the local node.

**localnodeid** (LONGINT): Node ID of the local node.

**numgroups** (LONGINT): Number of inconsistent groups.

## ctran.net events

### ctran.net.unreachable

### Severity

ALERT

### Description

This message occurs when the Cluster Transaction Manager (CTRAN) on the local node is unable to reach a peer node that it believes is in CLAM quorum.

## Corrective Action

Check all cluster interfaces on the node initiating this error message and the peer node specified in the message. \* Ensure all interfaces are in up state with "network interface show -role cluster". \* Ensure that each node initiating this message can reach their peer node with "cluster ping". \* Check the cluster interconnect switches and correct any issues found with the health or configuration found. If the error still persists KCS CTRAN logs will need to be collected for further analysis. Trigger an ALL autosupport by using, "autosupport invoke -node \* -message "ctran.net.unreachable" -type all", a diagnostic autosupport by using "autosupport invoke-diagnostic -node \* -message "ctran.net.unreachable" -subsystem kcs" and contact NetApp technical support support.

## Syslog Message

Node %s (id=%llu) is unable to reach peer node %s (id=%llu) on connection %s (ACN=0x%08llx).

## Parameters

**local\_node\_name** (STRING): Node name of the local node.

**local\_node\_id** (LONGINT): Node ID of the local node.

**peer\_node\_name** (STRING): Node name of the unreachable peer node.

**peer\_node\_id** (LONGINT): Node ID of the unreachable peer node.

**connection** (STRING): Name of connection that is unreachable.

**acn** (LONGINTEX): The ACN value.

# ctrl events

## ctrl.link events

### ctrl.link.degraded

#### Severity

ERROR

#### Description

This message occurs when the software detects either symbol errors or out-of-sequence errors on HA interconnect links.

#### Corrective Action

If possible, replace the degraded interconnect cables and reenale automatic switchover by using the "options -option-name rdma.link\_monitor -option-value on" command. If the problem persists after replacing both interconnect cables, replace the interconnect hardware on both nodes.

#### Syslog Message

Link errors detected on card in slot %d. Number of symbol errors (Port1:Port2): %d:%d; Number of OOS errors(RQ:SQ): %d:%d.

#### Parameters

**slot** (INT): Slot number of the HA Interconnect device.  
**sym\_err1** (INT): Symbol errors on the first port.  
**sym\_err2** (INT): Symbol errors on the second port.  
**rq\_oos** (INT): Out-of-sequence errors on the receive queue.  
**sq\_oos** (INT): Out-of-sequence errors on the send queue.

### ctrl.link.switchover

#### Severity

INFORMATIONAL

#### Description

This message occurs when the software initiates a switchover of the active interconnect link, which has been declared degraded.

#### Corrective Action

If possible, replace the degraded interconnect cable and reenale the automatic switchover by using the "options -option-name rdma.link\_monitor -option-value on" command. If the problem persists after replacing both the interconnect cables, replace the interconnect hardware on both nodes.

#### Syslog Message

Link switchover initiated to switch active interconnect link %s.

#### Parameters

**degraded\_port\_name** (STRING): Port name of the degraded link that is being switched from.

# ctrl.rdma.events

## ctrl.rdma.badSystemId

### Severity

NOTICE

### Description

This message occurs when an invalid system ID is found.

### Corrective Action

(None).

### Syslog Message

Invalid system ID %d. High-availability interconnect might not work properly.

### Parameters

**name** (INT): Invalid system ID value.

## ctrl.rdma.heartBeat

### Severity

INFORMATIONAL

### Description

This message logs heartbeat messages.

### Corrective Action

(None).

### Syslog Message

HA interconnect: %s.

### Parameters

**hb** (STRING): Heartbeat string.

## ctrl.rdma.partnerOK

### Severity

INFORMATIONAL

### Description

This message occurs when the partner resumes from pause state.

### Corrective Action

(None).

### Syslog Message

Partner node (system ID: %lld IP address: %s) has resumed from pause state.



## Parameters

**sysid** (LONGINT): Partner system ID.

**ipaddr** (STRING): Partner IP address.

## ctrl.rdma.partnerPause

### Severity

INFORMATIONAL

### Description

This message occurs when the partner enters pause state.

### Corrective Action

(None).

### Syslog Message

Partner node (system ID: %lld IP address: %s) has entered pause state.

## Parameters

**sysid** (LONGINT): Partner system ID.

**ipaddr** (STRING): Partner IP address.

## ctrl.rdma.portStateChange

### Severity

INFORMATIONAL

### Description

This message occurs when the state of the HA interconnect link changed frequently in the last hour.

### Corrective Action

Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabte the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. When one link is down on platforms with two available links, HA link redundancy is lost. However, HA services are still working and enabled on the secondary link.

### Syslog Message

HA interconnect: Link to %s changed state %d times in the last hour.

## Parameters

**ip** (STRING): IP address of the peer.

**num\_times** (INT): Number of times the state changed.

# daasync events

## daasync.capacity events

### daasync.capacity.changed

**Severity**

NOTICE

**Description**

This message occurs when the capacity data of a target device changes.

**Corrective Action**

(None).

**Syslog Message**

Capacity data has changed for device path %s with target address %s.

**Parameters**

**path** (STRING): Device path.

**targ\_addr** (STRING): Target address.

## daasync.inquiry events

### daasync.inquiry.changed

**Severity**

NOTICE

**Description**

This message occurs when the inquiry data of a target device changes.

**Corrective Action**

(None).

**Syslog Message**

Inquiry data has changed in path %s with target address %s.

**Parameters**

**path** (STRING): Device path.

**targ\_addr** (STRING): Target address.

## daasync.periph events

### daasync.periph.alloc.fail

**Severity**

ERROR

**Description**

This message occurs when the system cannot attach a drive.

**Corrective Action**

Use the "disk show -fields vmdisk-target-address" command to get the list of drives that are visible to the system and to verify that the missing drive is connected properly.

**Syslog Message**

Unable to attach new device %s due to status 0x%x.

**Parameters**

**device\_name** (STRING): Drive name.

**status** (INTEX): Status of the system.

# daregister events

## daregister.periph events

### daregister.periph.lost

#### Severity

ERROR

#### Description

This message occurs when a peripheral device is lost during registration.

#### Corrective Action

Check whether the device is connected correctly.

#### Syslog Message

Lost peripheral device during registration in path %s with respect to target address %s.

#### Parameters

**path** (STRING): Device path.

**targ\_addr** (STRING): Target address.

# dbfs events

## dbfs.api events

### dbfs.api

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) event related to an API operation occurs, such as an authentication failure.

#### Corrective Action

(None).

#### Syslog Message

A DBS API event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

#### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the API event. It displays as a string, in JSON format, or is empty.

## dbfs.assignments events

### dbfs.assignments

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) assignments event is reported, which occurs when the distribution of block data within the cluster changes. The DBS is responsible for managing the data that backs the FlexVols®.

#### Corrective Action

(None).

## Syslog Message

A DBS assignments event %s of type %s occurred for ServiceID %u on node %u/%s. DriveID = %u/%s. EventID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# db.block events

## db.block.repair.success

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) successfully completes a block service repair operation involving block rewrites to all BServices that are replicas for the block . See the previous data related events for further information. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

Contact NetApp technical support for data related events to assess the severity and for any required corrective actions.

## Syslog Message

A DBS block repair success event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbb.bulk events

## dbb.bulk.op

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to bulk volume operations such as backups, restores, making of Snapshot copies, or cloning occurs. If the original distributed block store event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is "0" if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbb.capacity events

## dbb.capacity.stranded

### Severity

NOTICE

### Description

This message occurs when the Distributed Block Store (DBS) detects that more than half of the cluster capacity is available on only one node. To achieve high availability (HA), the DBS effectively reduces the capacity of the largest node so that some of its capacity is stranded (unusable).

### Corrective Action

Adjust node disk capacities so that no node exceeds the capacity of the remaining nodes in the cluster.

### Syslog Message

Distributed Block Store has detected stranded capacity. Cluster fault type %s, fault ID %u, status %u. %s

## Parameters

**cfType** (STRING): Distributed Block Store cluster fault type of the associated object.

**cfID** (INT): Distributed Block Store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault: 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed Block Store cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

# db.clone events

## db.clone

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to volume cloning occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is "0" if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# db.cluster events

## db.cluster

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to cluster operations occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.



## Corrective Action

(None).

## Syslog Message

A DBS event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is "0" if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbs.cluster.iops.overprov

### Severity

ERROR

### Description

This message occurs when the distributed block store detects that the sum of all minimum quality of service (QoS) IOPS of flexible volumes is greater than the expected IOPS of the cluster. Minimum QoS cannot be maintained for all flexible volumes simultaneously. The distributed block store is responsible for managing the data that backs the flexible volumes.

### Corrective Action

Adjust QoS settings on one or more flexible volumes, or add more nodes to increase cluster IOPS capabilities.

### Syslog Message

The distributed block store detects that the sum of all minimum QoS IOPS is greater than the expected IOPS of the cluster. The cluster fault type is %s and fault id is %u.

### Parameters

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## dbs.cluster.master

## Severity

NOTICE

## Description

This message occurs when a Distributed Block Store (DBS) event related to cluster master operation occurs, such as cluster master being selected or cluster capacity limits being reached.

## Corrective Action

(None).

## Syslog Message

A DBS cluster master event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbS.db events

## dbS.db

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to its internal shared database occurs.

### Corrective Action

(None).

### Syslog Message

A DBS database event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.  
**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.  
**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.  
**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.  
**evtID** (INT): DBS event ID number.  
**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## **dbb.director events**

### **dbb.director.addrs.stable**

#### **Severity**

NOTICE

#### **Description**

This message occurs when data address assignments are stable. This means that all data address movement, graceful or ungraceful, is complete and remains stable until another event on the cluster requires subsequent data address movement.

#### **Corrective Action**

(None).

#### **Syslog Message**

Data address movement is complete and data address assignments are stable.

#### **Parameters**

(None).

### **dbb.director.aggrs.stable**

#### **Severity**

NOTICE

#### **Description**

This message occurs when aggregate assignments are stable. This means that all aggregate movement, graceful or ungraceful, is complete and remains stable until another event on the cluster requires subsequent aggregate movement.

#### **Corrective Action**

(None).

#### **Syslog Message**

Aggregate movement is complete and aggregate assignments are stable.

#### **Parameters**

(None).

# db.discon events

## db.discon.snapmirror.end

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) is unable to access snapmirror endpoint on the network. The DBS is responsible for managing the data that backs the FlexVols®.

### Corrective Action

Check network connectivity between the cluster and the remote snapmirror endpoint. Check the 1G management network.

### Syslog Message

The Distributed Block Store is unable to access the snapmirror endpoint from cluster. The cluster fault type is %s and fault id is %u.

### Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.disconn events

## db.disconn.cluster.pair

### Severity

ERROR

### Description

This message occurs when the distributed block store detects that a cluster pair is disconnected or configured incorrectly. The distributed block store is responsible for managing the data that backs the flexible volumes.

### Corrective Action

Check network connectivity between clusters. One of the clusters in the pair might have become misconfigured or disconnected. Remove the local pairing and retry pairing the clusters.

### Syslog Message

The distributed block store detects that the cluster pair is disconnected or configured incorrectly. The cluster fault type is %s and fault id is %u.

### Parameters

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## dbb.disconn.remote.node

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) detects that a remote node is not connected to cluster network. The DBS is responsible for managing the data that backs the FlexVols®.

### Corrective Action

Ping the remote nodes using jumbo frames to test network connectivity.

### Syslog Message

The Distributed Block Store has detected that a remote node is not connected to cluster network. The cluster fault type is %s and fault id is %u.

### Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## dbb.drive events

### dbb.drive

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to drive operations occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is "0" if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## **db.drives.capacity.mismatch**

### **Severity**

ALERT

### **Description**

This message occurs when Distributed Block Store (DBS) detects a drive capacity mismatch.

### **Corrective Action**

Check capacity of drive against other DBS drives on the node. All drives are expected to be the same size.

### **Syslog Message**

DBS has detected a drive capacity mismatch. Cluster fault node %u, node UUID %s, drives %s, drive UUIDs %s, type %s, fault ID %u, status %u. %s

### **Parameters**

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfDriveIDs** (STRING): IDs of the affected drives.

**cfDriveUuids** (STRING): UUIDs of the affected drives.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the hardware mismatch.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## **db.drives.events**

### **db.drives.failed**

#### **Severity**

ERROR

#### **Description**

This message occurs when Distributed Block Store (DBS) has one or more failed drives. This can be happen due to: no access, too many failures, a missing drive, inaccessible master service for the node, or the drive is locked and cannot be unlocked or the authentication key cannot be accessed.

## Corrective Action

Check network connectivity for the node. Replace the drives. Make sure that the authentication key is available.

## Syslog Message

DBS has detected a failed drive(s). Cluster fault node %u, node UUID %s, drives %s, drive UUIDs %s, type %s, cluster fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfDriveIDs** (STRING): IDs of the failed drives

**cfDriveUuids** (STRING): UUIDs of the failed drives

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# dbm.ekm events

## dbm.ekm.cert.alert

### Severity

ALERT

### Description

This message occurs when Distributed Block Store (DBS) detects that an External Key Management (EKM) key server configuration contains a certificate that will expire in less than 7 days.

### Corrective Action

Renew each listed certificate and update the associated EKM key server configuration.

### Syslog Message

DBS has detected that an EKM key server certificate is nearing expiration. Cluster fault type %s, fault ID %u, status %u. %s

### Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the certificates that are nearing expiration.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## dbm.ekm.cert.error

## Severity

ERROR

## Description

This message occurs when Distributed Block Store (DBS) detects that an External Key Management (EKM) key server configuration contains a certificate that will expire in less than 30 days.

## Corrective Action

Renew each listed certificate and update the associated EKM key server configuration.

## Syslog Message

DBS has detected that an EKM key server certificate is nearing expiration. Cluster fault type %s, fault ID %u, status %u. %s

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the certificates that are nearing expiration.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

# dbt.encl events

## dbt.encl.at.rest

### Severity

NOTICE

### Description

This message occurs when Distributed Block Store (DBS) enables or disables encryption at rest on a self-encrypting drive. Encrypting drives automatically encrypt and decrypt the data as it is written or read from the drive media. Enabling encryption at rest protects the data from unauthorized access if the drive is power-cycled. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS encryption at rest event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays



zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## db.ensemble events

### db.ensemble

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) event related to the database ensemble occurs, such as the inability to reach the database in a node.

#### Corrective Action

(None).

#### Syslog Message

A DBS ensemble event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

#### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

### db.ensemble.alert

#### Severity

ALERT

#### Description

This message occurs when the Distributed Block Store (DBS) detects the database connection to less than half of the ensemble nodes (of 3 or 5 total) is not available. The database still has quorum and is operational.

#### Corrective Action

Verify power for all nodes and their network connectivity. Wait 10 minutes and check that the fault status has changed to Resolved. If the condition persists, contact NetApp technical support.

## Syslog Message

DBS has lost network connectivity or power an ensemble node. Cluster fault type %s, fault ID %u, status %u. %s

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the degraded ensemble.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.exception events

## db.exception

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) is unable to process a cluster fault.

### Corrective Action

Wait 60 minutes and check that the fault status has changed to Resolved. If the condition persists, contact NetApp technical support.

## Syslog Message

DBS has detected a software exception. Cluster fault type %s, fault ID %u, status %u. %s

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the exception.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.fault events

## db.fault.checker

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to the fault checker occurs, such as excessive time for the periodic fault checker to execute.

## Corrective Action

(None).

## Syslog Message

A DBS fault checker event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbfs.file events

## dbfs.file.system.capacity.low

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) detects the mounted container file system ("/", "/persist", "/var/log") has capacity below 5 percent.

### Corrective Action

Add additional filesystem capacity to the node or remove any unneeded files until the fault is cleared.

### Syslog Message

DBS has detected low filesystem free space on node %u/%s. Cluster fault type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Universal unique identifier (UUID) of the associated node. If there is no associated node, zeroes display.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## dbfs.file.system.read.only

### Severity

ALERT

### Description

This message occurs when Distributed Block Store (DBS) detects a mounted container file system ("/, "/persist", "/var/log") that should be writeable is marked as read-only.

### Corrective Action

Verify required node file systems ("/, "/persist", "/var/log") are writeable. If not, resolve this issue. Contact NetApp technical support if cause cannot be found.

### Syslog Message

DBS has detected a file system is read-only. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the file system set to read-only.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## dbfs.generic.events

### dbfs.generic.cluster.fault

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) cluster fault is reported as a generic fault. The DBS is responsible for managing the data that backs the FlexVols®. Multiple faults are mapped to this one event.

### Corrective Action

None.

### Syslog Message

A DBS generic cluster fault occurred, fault %s (%u), sev %u, service %u, node %u/%s, drives %s.

### Parameters

**cfCodeName** (STRING): Name of the original DBS cluster fault code.

**cfCode** (INT): DBS cluster fault code number.

**cfSeverity** (INT): Severity of the original DBS cluster fault. This is different than the EMS severity.

**serviceID** (INT): Service ID that identifies the associated cluster service. It will be "0" if there is no

associated service.

**nodeID** (INT): Node ID number of the associated node. It will be "0" if there is no associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.

**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## db.generic.event

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event is reported as a generic event. The DBS is responsible for managing the data that backs the FlexVols®. Multiple DBS events are mapped to this one event. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

DBS generic event %s occurred, type %s, service %u, node %u/%s. drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## db.ie events

### db.ie

#### Severity

NOTICE

## Description

This message occurs when a Distributed Block Store (DBS) internal interface exception condition occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

## Corrective Action

(None).

## Syslog Message

A DBS IE event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# db.s.ikm events

## db.s.ikm

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event relating to Internal Key Management (IKM) occurs, such as purging a key.

### Corrective Action

(None).

### Syslog Message

DBS IKM event %s occurred, type %s, service %u, node %u/%s. drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It will be 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It will be 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier UUID of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It will be 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated

drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It may be a string, in JSON format, or empty.

## dbinstall events

### dbinstall

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) install event occurs, such as inability to install the system software on a pending node.

#### Corrective Action

(None).

#### Syslog Message

A DBS install event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

#### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbslimit events

### dbslimit

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) limit event occurs, such as reaching the recommended or maximum number of volumes or virtual volumes. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

#### Corrective Action

(None).

## Syslog Message

A DBS limit event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbm.maintenance events

## dbm.maintenance.mode

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) maintenance mode event occurs, such as that a node no longer exists. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

## Syslog Message

A DBS maintenance mode event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.



# dbms.master events

## dbms.master.didnt.start.cm

### Severity

NOTICE

### Description

This message occurs when the node fails to execute cluster master startup. The system attempts to start the cluster master again automatically.

### Corrective Action

(None).

### Syslog Message

Cluster master startup on service ID %u on node ID %u (UUID: %s) failed. Reason: %s.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**reason** (STRING): Reason cluster master startup failed.

## dbms.master.started.cm

### Severity

NOTICE

### Description

This message occurs when the node completes cluster master startup successfully.

### Corrective Action

(None).

### Syslog Message

Cluster master startup on service ID %u on node ID %u (UUID: %s) completed successfully.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node universal unique identifier (UUID) of the associated node.

# dbms.mem events

## dbms.mem.threshold.alert

### Severity

ALERT

## Description

This message occurs when Distributed Block Store (DBS) detects a node's container has very low memory.

## Corrective Action

Purge unused volumes and allow up to an hour for garbage collection between the Block and Slice services to run. Alternatively, add more nodes.

## Syslog Message

DBS has detected free memory is low on a node. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the memory usage threshold.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## db.mem.threshold.error

### Severity

ERROR

### Description

This message occurs when Distributed Block Store (DBS) detects a node's container has low memory.

### Corrective Action

Purge unused volumes and allow up to an hour for garbage collection between the Block and Slice services to run. Alternatively, add more nodes.

### Syslog Message

DBS has detected free memory is low on a node. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the memory usage threshold.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# dbns.network events

## dbns.network.error.fault

### Severity

ERROR

### Description

This message occurs when a distributed block store detects network configuration or connectivity errors on nodes in the cluster. Multiple error conditions are mapped to this one event. The distributed block store is responsible for managing the data that backs the flexible volumes.

### Corrective Action

Check and rectify network configuration parameters for configuration faults. For network connectivity faults, check physical components such as network interface card (NIC) port, switch port and ethernet cable. Replace the faulty component.

### Syslog Message

The distributed block store detects a network configuration cluster fault on node %u/%s. The cluster fault type is %s and fault ID is %u.

### Parameters

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Universal unique identifier (UUID) of the associated node. If there is no associated node, zeroes display.

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## dbns.network.event

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event related to storage cluster networking occurs.

### Corrective Action

Contact NetApp technical support for events related to network errors to assess the severity and for any required corrective actions.

### Syslog Message

A DBS software networking event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is "0" if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is "0" if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## db.network.mtu.check

### Severity

ERROR

### Description

This message occurs when a distributed block store fails to send messages between the storage cluster nodes using configured maximum transmission unit (MTU size). The distributed block store is responsible for managing the data that backs the flexible volumes.

### Corrective Action

Check and rectify network MTU size on cluster nodes and switches connected to it. Contact NetApp technical support.

### Syslog Message

The distributed block store detects a network MTU check cluster fault on node %u/%s. The cluster fault type is %s and fault ID is %u.

## Parameters

**nodeID** (INT): Node ID number of the associated node. It is "0" if there is no associated node.

**nodeUuid** (STRING): Universal unique identifier (UUID) of the associated node. If there is no associated node, zeroes display.

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## db.node events

### db.node.maint.mode.alert

### Severity

ALERT

## Description

This message occurs when Distributed Block Store (DBS) detects a node with a maintenance mode failure.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

DBS has detected a node with a maintenance mode failure. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the node in maintenance mode.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## **dbs.node.maint.mode.error**

### Severity

ERROR

### Description

This message occurs when Distributed Block Store (DBS) detects a node is in maintenance mode.

### Corrective Action

Disable maintenance mode when maintenance is complete.

### Syslog Message

DBS has detected a node is in maintenance mode. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the node in maintenance mode.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## **dbs.node.offline**

### Severity

ALERT

## Description

This message occurs when Distributed Block Store (DBS) detects the Cluster Master cannot communicate with the master service on a node.

## Corrective Action

Check for network connectivity issues and hardware errors.

## Syslog Message

DBS has detected the Cluster Master cannot communicate with the master service on a node. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the node offline.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.platform\_hardware events

## db.platform\_hardware

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) platform hardware event occurs, such as a drive hardware fault. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS platform hardware event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbms.prestartup events

### dbms.prestartup

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) prestartup event occurs. Several events might be posted during startup.

#### Corrective Action

(None).

#### Syslog Message

A DBS prestartup event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

#### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbms.provisioned events

### dbms.provisioned.space.full

#### Severity

ERROR

#### Description

This message occurs when the distributed block store detects that the overall provisioned capacity of the cluster has exceeded the slice reserve space threshold. The slice reserve space is only used when slices get reassigned due to slice drive failure. Therefore, the slice files can no longer be migrated during drive failure until this fault is resolved. The distributed block store is responsible for managing the data that backs the flexible volumes.

## Corrective Action

Add more provisioned space, or delete and purge volumes to resolve this fault.

## Syslog Message

The distributed block store detected that the overall provisioned capacity of the cluster is full. The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

# db.raid events

## db.raid.group.degraded

### Severity

ERROR

### Description

This message occurs when distributed block store detects that a RAID group in the cluster is degraded. The RAID group can continue to serve data, but performance might be adversely affected.

### Corrective Action

Look for disk error EMS messages, or use REST, Kubernetes, or kubectl calls to check for disk errors. If the errors cannot be resolved, contact NetApp technical support.

### Syslog Message

Distributed block store has detected a degraded RAID group. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node universal unique identifier (UUID) string of the associated node.

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the degraded RAID group.

**cfExtSrc** (STRING): Distributed block store cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## db.raid.group.not.oper

### Severity

ALERT



## Description

This message occurs when distributed block store detects that a RAID group is not operational and is unable to serve data.

## Corrective Action

Look for disk error EMS messages, or use REST, Kubernetes, or kubectl calls to check for disk errors. If the errors cannot be resolved, contact NetApp technical support.

## Syslog Message

Distributed block store has detected a nonoperational RAID group. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node universal unique identifier (UUID) string of the associated node.

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the RAID group that is not operational.

**cfExtSrc** (STRING): Distributed block store cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

# db.remote events

## db.remote.cluster

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event such as a change in cluster pair connectivity status occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS remote cluster event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated

drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## **db.remote.rep.cluster.full**

### **Severity**

ALERT

### **Description**

This message occurs when the distributed block store detects that the volumes have paused remote replication because the target storage cluster is full. The distributed block store is responsible for managing the data that backs the flexible volumes.

### **Corrective Action**

Free target space.

### **Syslog Message**

The distributed block store detects that the volumes have paused remote replication because the target storage cluster is full. The cluster fault type is %s and fault id is %u.

### **Parameters**

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## **db.remrep events**

### **db.remrep.async.dly.exceed**

#### **Severity**

ERROR

#### **Description**

This message occurs when the Distributed Block Store (DBS) is executing remote replication for a pair of volumes, but that it has not reached active state for 6 hours. The DBS is responsible for managing the data that backs the FlexVols®.

#### **Corrective Action**

Check network connectivity between clusters. Inspect slice service logs to see if some issue is preventing replication from continuing.

#### **Syslog Message**

The Distributed Block Store is attempting to perform remote replication that has not reached active state for 6 hours. The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## dbfs.remrep.snap.cluster.full

### Severity

ALERT

### Description

This message occurs when the Distributed Block Store (DBS) detects that Remote Replication of Snapshots is paused for associated volumes because target cluster is full. The DBS is responsible for managing the data that backs the FlexVols®.

### Corrective Action

Free space on the target volume.

### Syslog Message

The Distributed Block Store detected that Remote Replication of Snapshots is paused for associated volumes because target cluster is full. The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## dbfs.remrep.snapshots.exceed

### Severity

ALERT

### Description

This message occurs when the Distributed Block Store (DBS) detects that Remote Replication of Snapshots is paused for associated volumes because target volume has exceeded its snapshot limit. The DBS is responsible for managing the data that backs the FlexVols®.

### Corrective Action

Delete snapshots on the target volume.

### Syslog Message

The Distributed Block Store detected that Remote Replication of Snapshots is paused for associated volumes because target volume has exceeded its snapshot limit. The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.schedul.events

## db.schedul.e.action.error

### Severity

ERROR

### Description

This message occurs when the Distributed Block Store (DBS) is running one or more of the scheduled activities but the activity failed, for example running a scheduled create snapshot fails to complete. The fault clears if the scheduled activity runs again and succeeds, if the scheduled activity is deleted, or if the activity is paused and resumed. The DBS is responsible for managing the data that backs the FlexVols®.

### Corrective Action

Check the scheduler entry for issues.

### Syslog Message

The Distributed Block Store is attempting one or more scheduled activities which fails to complete. The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.s.secondary.events

## db.s.secondary.cache.thresh

### Severity

ALERT

### Description

This message occurs when the Distributed Block Store (DBS) detects that Slice volume secondary write cache has reached the first fullness warning threshold. Secondary cache starts to fill when Slice service write requests to the Block service (and by extension the FireStorm service) are not receiving replies. Client write performance may be reduced if this condition persists. The DBS is responsible for managing the data that backs the FlexVols®.

## Corrective Action

Verify that all nodes are online. Verify that all Block and FireStorm services, and associated aggregates, are online. Attempt repairs needed to bring nodes and aggregates online. If the reason for this condition cannot be found, contact NetApp technical support.

## Syslog Message

Slice secondary cache fullness threshold reached for service ID %u on node %u.

## Parameters

**serviceID** (INT): Service ID that identifies the associated cluster service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**cfDriveIDs** (STRING): List of the drive IDs associated with the fault. The list might be empty.

**cfDriveUuids** (STRING): List of the drive UUIDs associated with the fault. The list might be empty.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db.service events

## db.service

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) service event occurs, such as inability to start the FireStorm service.

### Corrective Action

(None).

### Syslog Message

A DBS service event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## db.service.not.running

### Severity

ALERT

### Description

This message occurs when Distributed Block Store (DBS) detects a service is not running.

### Corrective Action

Verify that the service is started within 10 minutes.

### Syslog Message

DBS has detected a service is not running. Cluster fault service ID %u, node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**serviceID** (INT): Service ID that identifies the associated cluster service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## db.slice events

### db.slice.operation

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) event relating to a slice service operation such as removing a metadata drive, slice reassignment i.e. balancing volumes, moving primaries, snapshot success, failure, expiration, group snapshot success, failure occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS slice event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## **dbs.slice.service.unhealthy**

### **Severity**

ALERT

### **Description**

This message occurs when the Distributed Block Store (DBS) is trying to migrate data away from an unresponsive Slice Service. The DBS is responsible for managing the data that backs the FlexVols®.

### **Corrective Action**

Expect DBS to automatically resolve this failure. The cluster is automatically decommissioning data and re-replicating its data onto other healthy drives.

### **Syslog Message**

DBS is migrating data away from unresponsive Slice service %u, node %u.

### **Parameters**

**serviceID** (INT): Service ID that identifies the associated DBS service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node. It will be zeroes if there is no associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

## **dbs.snapshot events**

### **dbs.snapshot.scheduler**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a Distributed Block Store (DBS) event related to scheduling snapshots occurs. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

## Corrective Action

(None).

## Syslog Message

A DBS scheduler event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbnsnmp events

## dbnsnmp.trap

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) reports an SNMP trap. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

A DBS SNMP trap event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.



# dbsoftware events

## dbsoftware.version.mismatch

### Severity

ALERT

### Description

This message occurs when Distributed Block Store (DBS) detects a software version mismatch in the cluster.

### Corrective Action

Bring the node out of maintenance mode. Check for cluster faults and upgrade again if the faults are cleared.

### Syslog Message

DBS has detected a software upgrade that has failed. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

### Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the software version that did not upgrade properly.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# dbssl events

## dbssl.node.cert.expire

### Severity

NOTICE

### Description

This message occurs when Distributed Block Store (DBS) detects that a Secure Socket Layer (SSL) certificate on a node is nearing expiration. No action is needed. The system resolves this issue by automatically updating the certificate.

### Corrective Action

(None).

### Syslog Message

DBS has detected that an SSL certificate is nearing expiration. Cluster fault node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node universal unique identifier (UUID) string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the certificates that are nearing expiration.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

## dbstat events

### dbstat

#### Severity

NOTICE

#### Description

This message occurs when a Distributed Block Store (DBS) statistics event occurs, such as inconsistent statistics. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

#### Corrective Action

(None).

#### Syslog Message

A DBS statistics event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbsw events

### dbsw.encr.at.rest

#### Severity

NOTICE

## Description

This message occurs when a Distributed Block Store (DBS) event relating to Software Encryption At Rest occurs, such as master key rekey completes or fails, rewrapping drive keys with new SEAR master key takes too long occur. Software Encryption At Rest when enabled, encrypts all data written, and decrypts all data read automatically in the software. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

## Corrective Action

(None).

## Syslog Message

A DBS SW Encryption at Rest event occurred, event %s, type %s, service %u, node %u/%s, drive %u/%s, eventID %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# db.unexpected events

## db.unexpected.exception

### Severity

NOTICE

### Description

This message occurs when a Distributed Block Store (DBS) exception occurs unexpectedly, such as the inability to analyze existing error information. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

### Corrective Action

(None).

### Syslog Message

An unexpected DBS exception %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

### Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.  
**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.  
**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.  
**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.  
**evtID** (INT): DBS event ID number.  
**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

## dbns.unresponsive events

### dbns.unresponsive.service.alrt

#### Severity

ALERT

#### Description

This message occurs when Distributed Block Store (DBS) Master Service detects an unresponsive service.

#### Corrective Action

Wait 10 minutes and check that the fault status has changed to Resolved. If the condition persists, contact NetApp technical support.

#### Syslog Message

DBS has detected an unresponsive service. Cluster fault service ID %u, node %u, node UUID %s, type %s, fault ID %u, status %u. %s

#### Parameters

**serviceID** (INT): Service ID that identifies the associated cluster service.  
**nodeID** (INT): Node ID number of the associated node.  
**nodeUuid** (STRING): Node UUID string of the associated node.  
**cfType** (STRING): DBS cluster fault type of the associated object.  
**cfID** (INT): DBS cluster fault ID number associated with the fault.  
**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).  
**cfDetails** (STRING): Name of the affected service.  
**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

### dbns.unresponsive.service.err

#### Severity

ERROR

#### Description

This message occurs when Distributed Block Store (DBS) detects an unresponsive service.

#### Corrective Action

Verify that the service is restarted within 10 minutes.

## Syslog Message

DBS has detected an unresponsive service. Cluster fault service ID %u, node %u, node UUID %s, type %s, fault ID %u, status %u. %s

## Parameters

**serviceID** (INT): Service ID that identifies the associated cluster service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID string of the associated node.

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Name of the affected service.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db upgrade events

## db upgrade

### Severity

ERROR

### Description

This message occurs when Distributed Block Store (DBS) detects a software upgrade is in progress longer than 24 hours.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

DBS has detected a software upgrade is in progress. Cluster fault type %s, fault ID %u, status %u. %s

## Parameters

**cfType** (STRING): DBS cluster fault type of the associated object.

**cfID** (INT): DBS cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Details of the upgrade.

**cfExtSrc** (STRING): DBS cluster fault externalSource label created by the CreateClusterFault API command and attached to the fault for testing purposes.

# db valence events

## db valence.test.bad

### Severity

NOTICE

## Description

This message occurs when a Distributed Block Store (DBS) valence test failure event occurs. Valence test are internal periodic tests. If the original DBS event specifies multiple drives, then an EMS event is generated for each drive.

## Corrective Action

(None).

## Syslog Message

A DBS valence test failure event %s of type %s occurred for Service ID %u on node %u/%s. Drive ID = %u/%s. Event ID = %u.

## Parameters

**evtMessage** (STRING): Description of the DBS event, including context details.

**evtType** (STRING): Type of the original DBS event.

**serviceID** (INT): Service ID that identifies the associated DBS service. It is 0 if there is no associated service.

**nodeID** (INT): Node ID number of the associated node. It is 0 if there is no associated node.

**nodeUuid** (STRING): Generated Universal Unique Identifier (UUID) of the associated node. It displays zeroes if there is no associated node.

**driveID** (INT): Drive ID number of the associated drive. It is 0 if there is no associated drive.

**driveUuid** (STRING): Generated UUID of the associated drive. It displays zeroes if there is no associated drive.

**evtID** (INT): DBS event ID number.

**evtDetails** (STRING): Specific details of the DBS event. It displays as a string, in JSON format, or is empty.

# dbb.volumes events

## dbb.volumes.degraded

### Severity

ALERT

### Description

This message occurs when the distributed block store detects that the listed volumes currently have only one copy of some of their data (they do not have a live secondary). This normally happens when a primary or live secondary slice service experiences remote procedure call (RPC) message timeouts to a peer slice service and changes it to a dead secondary. It can also happen briefly when a volume is created. The volume is current, syncing to a live secondary. The distributed block store is responsible for managing the data that backs the flexible volumes.

### Corrective Action

Check for network connectivity issues and hardware errors. There should be other faults if specific hardware components have failed. This fault clears when syncing completes and there is a live secondary slice service.

### Syslog Message

The distributed block store detects that the listed volumes currently have only one copy of some of their data (they do not have a live secondary). The cluster fault type is %s and fault id is %u.

## Parameters

**cfType** (STRING): Distributed block store cluster fault type of the associated object.

**cfID** (INT): Distributed block store cluster fault ID number associated with the fault.

**cfStatus** (INT): Current status of the cluster fault. 1 = New (just reported), 2 = Existing (updated), 3 = Resolved (closed).

**cfDetails** (STRING): Description of the cluster fault, including context details.

**cfExtSrc** (STRING): Distributed block store cluster fault external source label created by the "CreateClusterFault" API command and attached to the fault for testing purposes.

# dcx events

## dcx.framework events

### dcx.framework.disabled

**Severity**

NOTICE

**Description**

This message occurs when the DCS framework has been disabled on this node. This node will not allow any updates to any counters stored in the DCS framework on this node. However, the values for existing counters may be retrieved for display purposes.

**Corrective Action**

(None).

**Syslog Message**

The DCS framework is disabled on this node.

**Parameters**

(None).

### dcx.framework.enabled

**Severity**

INFORMATIONAL

**Description**

This message occurs when the DCS framework has been enabled on this node.

**Corrective Action**

(None).

**Syslog Message**

The DCS framework is enabled on this node.

**Parameters**

(None).



# dev events

## dev.driver events

### dev.driver.throttling

#### Severity

NOTICE

#### Description

This message occurs when an I/O driver is throttling I/O requests because the underlying controller is overloaded with I/O requests.

#### Corrective Action

(None).

#### Syslog Message

%s driver unit %d throttling I/O requests due to %s.

#### Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the controller.

**reason** (STRING): Reason for throttling.

**current\_throttle\_sendq** (INT): Current throttle value for send requests

**lowest\_throttle\_sendq** (INT): Lowest throttle value for send requests

## dev.path events

### dev.path.inquiry.failed

#### Severity

ERROR

#### Description

This message occurs when the system fails to discover the drive.

#### Corrective Action

Use the "disk show -fields vmdisk-target-address" command to get the list of drives that are visible to the system and to verify that the missing drive is connected properly.

#### Syslog Message

Failed to discover drive (device name:%s%d).

#### Parameters

**periph\_name** (STRING): Drive name.

**unit\_no** (INT): Drive unit number.

# device events

## device.install events

### device.install.successful

#### Severity

INFORMATIONAL

#### Description

This message occurs when Data ONTAP® successfully installs one or more devices after a prior installation failure.

#### Corrective Action

(None).

#### Syslog Message

Device type %s installation successful for one or more devices that reported installation failure previously.

#### Parameters

**devtype** (STRING): SCSI device type.

## device.latency events

### device.latency.threshold

#### Severity

NOTICE

#### Description

This message occurs when a device exceeds the latency threshold for the 30- second monitoring period. The event provides the name of the device, the target device (if applicable), the current threshold value, the average execution latency, and total latency for this period. The event also displays the number of I/O operations for the period, which is used to calculate latency.

#### Corrective Action

Contact your hypervisor provider if this issue persists. This usually indicates undersized storage in terms of allowed operations per second.

#### Syslog Message

The device latency in microseconds exceeded the threshold value for device "%s%d" and the associated target device "%s". The threshold is "%d", execution latency is "%d", total latency is "%d", and the number of operations for this period is "%d".

#### Parameters

**devname** (STRING): Name of the device. The device name can be "isl" (local to this node), "isp" (local to the partner node) or "ism" (the mediator with no associated target device), which represents the iSCSI name. Any other device name will be the name of the actual target device. If the device name is not "isl", or "isp", or the target device name, the name will always be "n/a".

**unit** (INT): Unit number of the device.

**target\_dev** (STRING): Name of the target device. The name will be the target (end point) device name if the device name is either "isl" or "isp". Otherwise it is the actual underlying device on this node.

**threshold** (INT): Monitoring threshold in microseconds.

**exec\_lat** (INT): Execution latency in microseconds for this period.

**total\_lat** (INT): Total latency in microseconds for this period.

**num\_operations** (INT): Number of I/O operations for this period.

## device.max events

### device.max.count.exceeded

#### Severity

ERROR

#### Description

This message occurs when Data ONTAP® cannot install one or more discovered devices since the device type count is at full capacity.

#### Corrective Action

If possible, remove an existing device of the same type so that device installation can proceed. If the failure persists, contact NetApp technical support.

#### Syslog Message

Unable to install one or more %s devices discovered on initiator port %s, target port WWNN: %s, WWPN: %s. The device limit (%d) has been reached.

#### Parameters

**devtype** (STRING): SCSI device type.

**adapter** (STRING): Name of the initiator port discovering the device(s).

**wwnn** (STRING): World Wide Node Name of the target port exporting the device(s).

**wwpn** (STRING): World Wide Port Name of the target port exporting the device(s).

**limit** (INT): Device maximum count.

# dfu events

## dfu.badbridgeshelfconfig events

### dfu.badBridgeShelfConfig

#### Severity

ERROR

#### Description

This message occurs when the disk firmware update aborts because it cannot contact the two peer FC-to-SAS bridges.

#### Corrective Action

Make sure that the storage controller is correctly cabled to the two peer FC-to-SAS bridges and that the bridges are correctly cabled to the SAS disk shelves. Then retry the 'disk\_fw\_update' command.

#### Syslog Message

Error contacting two peer FC-to-SAS bridges. Unable to download firmware to disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.clusterdisknotready events

### dfu.clusterDiskNotReady

#### Severity

ERROR

#### Description

This message occurs when the controller failover partner reports that a disk selected for download is in the NOT READY state.

#### Corrective Action

Only disks in the READY state can be downloaded to. If a disk is in the NOT READY state, it indicates that the disk has probably failed for some reason and should be returned for RMA.

#### Syslog Message

Controller failover partner reports disk as NOT READY. Unable to download firmware to disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.

**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.clustericdown events

### dfu.clustericDown

#### Severity

ERROR

#### Description

This message occurs when the high-availability (HA) interconnect is not available.

#### Corrective Action

The HA interconnect was not functional. Make sure that controller failover is enabled. Check that the HA interconnect cables are connected properly, and if applicable (based on the HA interconnect type), that the SFPs (Small Form-Factor Pluggables) are seated properly on both controllers.

#### Syslog Message

HA interconnect not available (cf disabled?). Unable to download firmware to disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.clustericerror events

### dfu.clustericError

#### Severity

ERROR

#### Description

This message occurs when communication with the controller failover partner cannot be established.

#### Corrective Action

Typically, these errors are related to cable or transmission errors from the HA interconnect. Check that the HA interconnect cables are connected properly, and if applicable (based on the HA interconnect type), that the SFPs (Small Form-Factor Pluggables) are seated properly on both controllers.

#### Syslog Message

HA interconnect error %d encountered. Unable to download firmware to disk %s [%s %s %s] S/N [%s].

#### Parameters

**clusterError** (INT): High-availability (HA) interconnect error number.  
**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.

**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.clusterpartner events

### dfu.clusterPartner

#### Severity

ERROR

#### Description

This message occurs when automatic firmware download is performed as part of system boot and partner disks are downrev.

#### Corrective Action

Firmware download is not allowed to update partner disks automatically. Manual invocation of the "storage disk firmware update" command is required. See the "storage disk firmware update" man page for details.

#### Syslog Message

Operator must issue disk\_fw\_update command to download firmware to partner Disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.clustertakeover events

### dfu.clusterTakeover

#### Severity

ERROR

#### Description

This message occurs when firmware download is attempted while in takeover mode.

#### Corrective Action

Firmware download is not allowed while in takeover mode. Exit takeover mode and attempt the "storage disk firmware update" command again. See the "storage disk firmware update" man page for details.

#### Syslog Message

Unable to download firmware while in takeover mode to Disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product id of the disk.

**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.disksizechanged events

### dfu.diskSizeChanged

#### Severity

NOTICE

#### Description

This message occurs when a size change is detected after a firmware update. This can be either an increase or decrease in physical size of the disk.

#### Corrective Action

The disk might not be usable any longer by the RAID subsystem because of the size change that occurred. If RAID cannot access the disk any longer, it will be failed and should be returned for RMA. If the disk is still accessible, leaving it in the system is perfectly fine.

#### Syslog Message

Disk %s [%s %s %s] S/N [%s] size changed (%x:%x %llx:%llx %llx:%llx %llx:%llx %llx:%llx)

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.  
**oldSectorSize** (INTHEX): Old sector size of the disk.  
**newSectorSize** (INTHEX): New sector size of the disk.  
**oldLastSector** (LONGINTEX): Old right-size of the disk.  
**newLastSector** (LONGINTEX): New right-size of the disk.  
**oldDiskBlocks** (LONGINTEX): Old right-sized blocks on the disk.  
**newDiskBlocks** (LONGINTEX): New right-sized blocks on the disk.  
**oldRealLastSector** (LONGINTEX): Old raw size of the disk.  
**newRealLastSector** (LONGINTEX): New raw size of the disk.  
**oldRealDiskBlocks** (LONGINTEX): Old raw blocks on the disk.  
**newRealDiskBlocks** (LONGINTEX): New raw blocks on the disk.

## dfu.firmwaredownloading events

### dfu.firmwareDownloading

#### Severity

INFORMATIONAL

#### Description

This message occurs when firmware is being downloaded for a specific product ID of disks on the system.

### Corrective Action

(None).

### Syslog Message

Now downloading firmware file %s on %d disk(s) of plex [%s]...

### Parameters

**file** (STRING): Firmware file which is being downloaded.

**diskcount** (INT): Number of disks being downloaded to.

**plex** (STRING): Plex name being downloaded to.

## dfu.firmwaredownrev events

### dfu.firmwareDownrev

#### Severity

ERROR

#### Description

This message occurs when the system finds disks with a downrev version of firmware, as defined by the installed version of Data ONTAP®.

#### Corrective Action

Use the "storage disk firmware update" command to update the downrev firmware on the disks. Check your support provider for more information about obtaining the latest disk firmware. If the problem persists, contact NetApp technical support.

#### Syslog Message

Downrev firmware on %d disks.

#### Parameters

**downrev** (INT): Number of disks with downrev firmware.

## dfu.firmwareuptodate events

### dfu.firmwareUpToDate

#### Severity

NOTICE

#### Description

This message occurs when all eligible disks are determined to be updated to the latest version of firmware available on the system.

#### Corrective Action

(None).

#### Syslog Message

Firmware is up-to-date on all eligible disks.



## Parameters

(None).

# dfu.fwdownloadfailedmultfiles events

## dfu.fwDownloadFailedMultFiles

### Severity

ERROR

### Description

This message occurs when more than one file is found for downloading to an ATA drive. Because there is no version string control, only one file at a time is allowed on the system.

### Corrective Action

Two or more firmware files were found for download. Delete all except the one that should be downloaded to the drive. Background firmware update should download on the next check cycle.

### Syslog Message

Download failed due to multiple firmware files found for Disk type [%s %s].

### Parameters

**vendorName** (STRING): Vendor name of the disk.

**productId** (STRING): Product ID of the disk.

# dfu.fwdownloadfailedstatus events

## dfu.fwDownloadFailedStatus

### Severity

ERROR

### Description

This message occurs when a firmware download failure occurs.

### Corrective Action

Decode sense codes based on drive supplier to determine the reason for the failure. Most situations can be recovered by attempting to download firmware again using the "disk\_fw\_update" command. See the disk\_fw\_update man page for complete details.

### Syslog Message

Status: (0x%x/0x%x/0x%x) (SCSI sense 0x%x/0x%x/0x%x/0x%x)

### Parameters

**OpStatus** (INTEX): Status of the passthru I/O.

**HostStatus** (INTEX): Host Status of the passthru I/O.

**TargetStatus** (INTEX): Target Status of the passthru I/O.

**sense\_key** (INTEX): SCSI Sense key.

**sense\_code** (INTEX): SCSI Sense Code.

**sense\_code\_qualifier** (INTEX): SCSI Sense Code Qualifier.

**fru** (INTHEX): SCSI Sense Code FRU identifier.

## dfu.fwdownloadfailedwbuf events

### dfu.fwDownloadFailedWbuf

#### Severity

ERROR

#### Description

This message occurs when retries are exhausted on the write buffer command that is used to download firmware to a disk.

#### Corrective Action

Disk firmware might still need to be updated. See the `disk_fw_update` man page for complete details on performing this.

#### Syslog Message

Write buffer command failed to Disk %s [%s %s %s] S/N [%s].

#### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.init events

### dfu.init.noMemory

#### Severity

ERROR

#### Description

This message occurs when the disk firmware update initialization routine cannot allocate memory.

#### Corrective Action

Memory allocation relies on the availability of memory resources in the system. If possible, reduce the workload so that fewer resources are needed. If the failure persists after reducing the workload, contact NetApp technical support.

#### Syslog Message

Could not allocate memory for disk firmware update. Requested size=%d bytes.

#### Parameters

**size** (INT): Size (in bytes) of the memory allocation request.

# dfu.mallocfailed events

## dfu.mallocFailed

### Severity

ALERT

### Description

This message occurs when Data ONTAP® cannot allocate memory to read the disk firmware file and aborts the disk firmware update.

### Corrective Action

Inability to allocate memory for this action indicates a failure within Data ONTAP. Take a core for reference, and then open a support case to address the underlying issue. The reboot process when you take the core resolves the immediate problem.

### Syslog Message

Cannot allocate memory to read disk firmware file %s; disk firmware update aborted.

### Parameters

**filename** (STRING): Name of the disk firmware file.

# dfu.nomemory events

## dfu.noMemory

### Severity

ERROR

### Description

This message occurs when the disk firmware update routine cannot allocate memory to update the disk firmware.

### Corrective Action

Memory allocation relies on the availability of memory resources in the system. If possible, reduce the workload so that fewer resources are needed. Rebooting the system can recover memory resources that might have been lost. If the failure persists after reducing the workload and rebooting, contact NetApp technical support for assistance in determining whether additional system memory is indicated.

### Syslog Message

Could not allocate memory for disk firmware update. Unable to download firmware to disk %s [%s %s %s] S/N [%s].

### Parameters

**diskName** (STRING): Name of the disk.

**vendorName** (STRING): Name of the vendor of the disk.

**productId** (STRING): Product ID of the disk.

**fwVersion** (STRING): Firmware version of the disk.

**serialno** (STRING): Serial number of the disk.

# dfu.readtestfailed events

## dfu.readTestFailed

### Severity

ERROR

### Description

This message occurs when a read of the disk fails after downloading firmware to a disk.

### Corrective Action

Disk was unable to be read from after a firmware update occurred. I/O will be retried. If the retries are not successful, you should fail the disk and return it for RMA.

### Syslog Message

Media Access test failed on Disk %s [%s %s %s] S/N [%s].

### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

## dfu.readTestFailed.fatal

### Severity

ERROR

### Description

This message occurs when a read of the disk fails after downloading firmware to a disk and all retries have been exhausted.

### Corrective Action

Disk was unable to be read from after a firmware update occurred. I/O will be retried. If the retries are not successful, you should fail the disk and return it for RMA.

### Syslog Message

Unable to access media after firmware update on Disk %s [%s %s %s] S/N [%s].

### Parameters

**diskName** (STRING): Name of the disk.  
**vendorName** (STRING): Vendor name of the disk.  
**productId** (STRING): Product ID of the disk.  
**fwVersion** (STRING): Firmware version of the disk.  
**serialno** (STRING): Serial number of the disk.

# dfu.readtestfailedstatus events

## dfu.readTestFailedStatus

### Severity

ERROR

### Description

This message occurs when a read test failure occurs after a firmware download.

### Corrective Action

Decode sense codes based on drive supplier to determine the reason for the failure. Most situations can be recovered by the system automatically retrying the I/O operation.

### Syslog Message

Status: (0x%x/0x%x/0x%x) (SCSI sense 0x%x/0x%x/0x%x/0x%x)

### Parameters

**OpStatus** (INTEHEX): Status of the read I/O.

**HostStatus** (INTEHEX): Host Status of the read I/O.

**TargetStatus** (INTEHEX): Target Status of the read I/O.

**sense\_key** (INTEHEX): SCSI Sense key.

**sense\_code** (INTEHEX): SCSI Sense Code.

**sense\_code\_qualifier** (INTEHEX): SCSI Sense Code Qualifier.

**fru** (INTEHEX): SCSI Sense Code FRU identifier.

## dfu.refreshinquiryfailed events

### dfu.refreshInquiryFailed

### Severity

ERROR

### Description

This message occurs after a successful disk firmware update, when the 'SCSI inquiry' command, which is issued to all disks as part of the disk firmware download process, fails to finish. It is also possible to get messages indicating that disks have downrev firmware.

### Corrective Action

When this error happens, check the connection to the partner node and manually retry the firmware download process. If the problem still persists, open a support case to investigate the underlying cause.

### Syslog Message

(None).

### Parameters

**retryCount** (INT): Number of times the command was retried and failed.

## dfu.suspended\_disasterrecovery events

## **dfu.suspended\_DisasterRecovery**

### **Severity**

ERROR

### **Description**

This message occurs in a MetroCluster™ environment when firmware download to remote disks is suspended during disaster recovery. This event is signaled at most once during disaster recovery, the first time firmware download to a remote disk is suspended.

### **Corrective Action**

Firmware download to remote disks is suspended during MetroCluster disaster recovery. Wait until switchback has completed, and either allow the firmware to be downloaded automatically, or run the "storage disk firmware update" command. See the "storage disk firmware update" man page for details.

### **Syslog Message**

Suspend firmware download to remote disks while in MetroCluster Disaster Recovery state %s

### **Parameters**

**dr\_state** (STRING): Current MetroCluster disaster recovery state.

# diff events

## diff.svcs events

### diff.svcs.setting.topartner

#### Severity

ERROR

#### Description

This message occurs when the system cannot automatically apply the Differentiated Services (diff-svcs) setting to the High Availability (HA) partner. In the event there is a setting mismatch between HA pairs the node with the incorrect setting may produce an incorrect At-Risk-System (ARS) report.

#### Corrective Action

Manually apply the diff-svcs setting to the HA partner to ensure the setting is consistent across both nodes in the HA pair. In the event of a failure, issue the following command on the HA partner: "system node modify -node nodename -diff-svcs true|false" .

#### Syslog Message

Differentiated services setting was not automatically applied to the HA partner of node %s: %s.

#### Parameters

**node** (STRING): Name of the node where the setting was successfully applied. A failure occurred while trying to apply the same setting to its HA partner.

**error** (STRING): Error reported when attempting to apply the setting to the HA partner.

# dipmon events

## dipmon.dataip events

### dipmon.dataip.responsive

#### Severity

NOTICE

#### Description

This message occurs when a data IP address becomes responsive on the node on which it is hosted and has been committed to the Zookeeper database.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s (ID: %u) hosted on node %s (ID: %u, Uuid: %s) is responsive and has been committed to the database.

#### Parameters

**dataip** (STRING): IP address of the data IP.

**dataipID** (INT): Unique identifier of the data IP.

**nodeName** (STRING): Name of the node that is hosting the data IP.

**nodeID** (INT): Unique identifier of the node that is hosting the data IP.

**nodeUuid** (STRING): Universal unique identifier of the node that is hosting the data IP.

### dipmon.dataip.unresponsive

#### Severity

NOTICE

#### Description

This message occurs when a data IP address becomes unresponsive on the node on which it is hosted and has been committed to the Zookeeper database.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s (ID: %u) hosted on node %s (ID: %u, Uuid: %s) is unresponsive and has been committed to the database.

#### Parameters

**dataip** (STRING): IP address of the data IP.

**dataipID** (INT): Unique identifier of the data IP.

**nodeName** (STRING): Name of the node that is hosting the data IP.

**nodeID** (INT): Unique identifier of the node that is hosting the data IP.

**nodeUuid** (STRING): Universal unique identifier of the node that is hosting the data IP.



# director events

## director.aggr events

### director.aggr.reassigned

#### Severity

NOTICE

#### Description

This message occurs when an aggregate is being reassigned from one node to another node. The output contains the Firefly ID and node ID of both nodes as well as whether or not the operation was done forcefully. The value of isForceful is "true" if an aggregate is moved in response to a error condition on the cluster rather than a user requested operation.

#### Corrective Action

(None).

#### Syslog Message

Aggregate (ID: %u) is being reassigned (forcefully:%u) from node %s (ID: %u, Uuid: %s Firefly ID: %u) to node %s (ID: %u, Uuid: %s Firefly ID: %u).

#### Parameters

**aggrID** (INT): Unique identifier of the aggregate being reassigned.

**isForceful** (INT): Boolean value to denote whether or not the reassignment was done forcefully.

**fromNode** (STRING): Name of the node that the aggregate is being reassigned from.

**fromNodeID** (INT): Unique identifier of the node that the aggregate is being reassigned from.

**fromNodeUuid** (STRING): Universal unique identifier of the node that the aggregate is being reassigned from.

**fromFireflyID** (INT): Unique identifier of the Firefly Service that the aggregate is reassigned to.

**toNode** (STRING): Name of the node that the aggregate is being reassigned to.

**toNodeID** (INT): Unique identifier of the node that the aggregate is being reassigned to.

**toNodeUuid** (STRING): Universal unique identifier of the node that the aggregate is being reassigned to.

**toFireflyID** (INT): Unique identifier of the Firefly Service that the aggregate is reassigned to.

## director.dataip events

### director.dataip.reassigned

#### Severity

NOTICE

#### Description

This message occurs when a data IP address is being reassigned from one node to another node. The output contains the node ID of both nodes as well as whether or not the operation was done forcefully. The value of isForceful is "true" if the data IP is moved in response to a error condition on the cluster rather than a user requested operation.

## Corrective Action

(None).

## Syslog Message

Data IP address %s (ID: %u) is being reassigned (forcefully:%u) from node %s (ID: %u, Uuid: %s) to node %s (ID: %u, Uuid: %s).

## Parameters

**dataip** (STRING): IP address of the data IP.

**dataipID** (INT): Unique identifier of the data IP.

**isForceful** (INT): Boolean value to denote whether or not the reassignment was done forcefully.

**fromNode** (STRING): Name of the node from which the data IP is being reassigned.

**fromNodeID** (INT): Unique identifier of the node from which the data IP is being reassigned.

**fromNodeUuid** (STRING): Universal unique identifier of the node from which the data IP is being reassigned.

**toNode** (STRING): Name of the node to which the data IP is being reassigned.

**toNodeID** (INT): Unique identifier of the node to which the data IP is being reassigned.

**toNodeUuid** (STRING): Universal unique identifier of the node to which the data IP is being reassigned.

# discover events

## discover.6500 events

### discover.6500.supported

**Severity**

NOTICE

**Description**

This message occurs when discovery of a 6500N FC-to-SAS bridge is enabled.

**Corrective Action**

(None).

**Syslog Message**

FC-to-SAS bridge ATTO 6500N discovery is enabled.

**Parameters**

(None).

### discover.6500.unsupported

**Severity**

NOTICE

**Description**

This message occurs when discovery of a 6500N FC-to-SAS bridge is disabled.

**Corrective Action**

(None).

**Syslog Message**

FC-to-SAS bridge ATTO 6500N discovery is disabled.

**Parameters**

(None).

# discoveryd events

## discoveryd.search events

### discoveryd.search.request

**Severity**

NOTICE

**Description**

This message occurs when the discoveryd service is contacted during the initial cluster setup phase with an M-SEARCH multicast and a follow up HTTP GET to get the device description document.

**Corrective Action**

(None).

**Syslog Message**

The discoveryd service was contacted by a device (IP address: %s) to discover information about this node.

**Parameters**

**client\_ip** (STRING): IP address of the client machine requesting the discoveryd information.

# disk events

## disk.adaptererror events

### disk.adapterError

**Severity**

ALERT

**Description**

This message occurs when an adapter error is detected.

**Corrective Action**

(None).

**Syslog Message**

Disk %s: unexpected adapter error 0x%x

**Parameters**

**diskName** (STRING): Name of the device.

**error** (INHEX): Adapter status.

## disk.bdfu events

### disk.bdfu.suspended

**Severity**

ERROR

**Description**

This message occurs when the background drive firmware update process fails to unlock SAS NetApp® Storage Encryption (NSE) drives in FIPS-compliance mode. The FIPS-certified SAS NSE drives must be unlocked with a valid authentication key before firmware can be downloaded to them. If the number of failed unlock attempts reaches a certain threshold, the drives will be locked out, leading to permanent loss of data in certain conditions. The system will automatically suspend the background drive firmware update for 12 hours, after which it is reenabled.

**Corrective Action**

Contact NetApp technical support for assistance with restoring valid authentication keys on the drive.

**Syslog Message**

Background drive firmware update is suspended because FIPS NSE drives cannot be unlocked. Reason: %s.

**Parameters**

**suspend\_reason** (STRING): Reason for the drive unlock failure.

# disk.capacity events

## disk.capacity.changed

### Severity

NOTICE

### Description

This message occurs when a media size change is reported for the underlying storage device. The new size and previous size are measured in 4096-byte blocks.

### Corrective Action

(None).

### Syslog Message

A storage device reported a size change for device "%s". New size: %llu, previous size: %llu (4096-byte blocks).

### Parameters

**diskName** (STRING): Name of the disk.

**new\_size** (LONGINT): The new size measured in 4096-byte blocks.

**previous\_size** (LONGINT): The previous size measured in 4096-byte blocks.

# disk.checksum events

## disk.checksum.offlineAdapter

### Severity

ALERT

### Description

This message occurs when multiple errors are detected during checksum verification. The adapter is taken offline.

### Corrective Action

Replace the adapter. Contact Contact NetApp technical support. for assistance with a Return Material Authorization.

### Syslog Message

Adapter %s taken offline due to checksum verification failure on multiple disks. Keep the adapter offline and contact Contact NetApp technical support. for assistance.

### Parameters

**adapterName** (STRING): Name of the adapter.

## disk.checksum.verifyFailed

### Severity

ALERT

## Description

This message occurs when disk checksum verify fails.

## Corrective Action

(None).

## Syslog Message

Disk level checksum verification failed on WRITE VERIFY: Disk %s, Block #%llu: Volume %s, FileId %d, Block #%llu

## Parameters

**diskName** (STRING): Name of the device.

**bno** (LONGINT): block id

**vol** (STRING): Volume name.

**fileid** (INT): File ID.

**block** (LONGINT): File block ID.

# disk.cksum events

## disk.cksum.bno.verifyFailed

### Severity

ALERT

### Description

This message occurs when the system detects a block number mismatch during disk checksum verification. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location.

### Corrective Action

(None).

### Syslog Message

Block number mismatch on WRITE VERIFY: Disk %s, expected\_dbn = %u, stored\_dbn = %u, expected\_vbn = %llu, stored\_vbn = %llu, Volume %s, FileId %d, Block #%llu.

### Parameters

**diskName** (STRING): Name of the device.

**expected\_dbn** (INT): Expected physical disk block number.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected volume block number.

**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.

**vol** (STRING): Volume name.

**fileid** (INT): File ID.

**block** (LONGINT): File block ID.

# disk.ddr events

## **disk.ldr.dqp.out.of.date**

### **Severity**

NOTICE

### **Description**

This message occurs when the Dynamic Drive Recognition (DDR) Scanner detects a disk that is not part of its qualification table.

### **Corrective Action**

Check your support provider's web site or knowledgebase for more information about downloading and installing an updated Disk Qualification Package (DQP) file.

### **Syslog Message**

The DDR Scanner has detected that the latest DQP is not present on this system.

### **Parameters**

(None).

## **disk.ldr.invalid.azcs.capacity**

### **Severity**

ERROR

### **Description**

This event is generated when the DDR Scanner detects a capacity value that does not match with the alias string associated with it.

### **Corrective Action**

(None).

### **Syslog Message**

The DDR Scanner has detected an invalid AZCS alias capacity on disk %s

### **Parameters**

**diskname** (STRING): Name of the disk.

**capacity\_value** (LONGINT): The value of the invalid capacity

## **disk.ldr.invalid.bcs.capacity**

### **Severity**

ERROR

### **Description**

This event is generated when the DDR Scanner detects a capacity value that does not match with the alias string associated with it.

### **Corrective Action**

(None).



## Syslog Message

The DDR Scanner has detected an invalid BCS alias capacity on disk %s

## Parameters

**diskname** (STRING): Name of the disk.

**capacity\_value** (LONGINT): The value of the invalid capacity

## disk.ddr.scanner.ddr.failure

### Severity

NOTICE

### Description

This event is issued when a the ddr scanner has failed to read the new DDR label it just placed on the disk.

### Corrective Action

(None).

## Syslog Message

Disk %s Shelf %s Bay %s [%s %s %s] S/N [%s] [%s]: DDR label verify Failure - DDR Error

## Parameters

**disk\_name** (STRING): The name of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision of the disk

**serialno** (STRING): Serial number of the disk

**ddr\_error** (STRING): DDR Error string which has been parsed

## disk.ddr.scanner.label.done

### Severity

NOTICE

### Description

This event is issued when a the ddr scanner has completed has completed scanning a disk successfully

### Corrective Action

(None).

## Syslog Message

Disk %s Shelf %s Bay %s [%s %s %s] S/N [%s] : DDR label write complete

## Parameters

**disk\_name** (STRING): The name of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision of the disk  
**serialno** (STRING): Serial number of the disk

## disk.ddr.scanner.verify.failure

### Severity

NOTICE

### Description

This event is issued when a the ddr scanner has failed to read the new DDR label it just placed on the disk.

### Corrective Action

(None).

### Syslog Message

Disk %s Shelf %s Bay %s [%s %s %s] S/N [%s] : DDR label verify Failure

### Parameters

**disk\_name** (STRING): The name of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision of the disk  
**serialno** (STRING): Serial number of the disk

## disk.ddr.scanner.write.failure

### Severity

NOTICE

### Description

This event is issued when a the ddr scanner has failed to write a new DDR label to the disk

### Corrective Action

(None).

### Syslog Message

Disk %s Shelf %s Bay %s [%s %s %s] S/N [%s] : DDR label write Failure

### Parameters

**disk\_name** (STRING): The name of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision of the disk  
**serialno** (STRING): Serial number of the disk

## disk.ldr.unable.add.disk

### Severity

ALERT

### Description

This event is generated when the DDR Scanner is unable to re-add a disk to the filer after trying for 2 minutes. This means there are no more spare disk\_Cbs in the filer.

### Corrective Action

(None).

### Syslog Message

The filer has not freed a disk\_Cb for over 2 minutes for %s.

### Parameters

**diskuid** (STRING): The unique identifier of the disk

## disk.diagnostic events

### disk.diagnostic.logs.saved

### Severity

NOTICE

### Description

This message occurs when ONTAP retrieves the diagnostic logs from the disk. The logs are saved at `/etc/logs/drive_failure_logs/`.

### Corrective Action

(None).

### Syslog Message

Diagnostic logs retrieved for disk "%s".

### Parameters

**diskName** (STRING): Name of the disk.

## disk.duplicate events

### disk.duplicate.name

### Severity

ERROR

### Description

This message occurs when duplicate disk names are detected. This can be caused by corruption of the disk's shelf id or bay number or failure to determine the stack id. There is no known impact to functional systems, but failure to correct this error may result in operations taken on the wrong disk and ambiguity where the disks are referenced.

## Corrective Action

If any disk reports a shelf id or bay number not corresponding to its physical location, choose one of the following options: 1. Perform a takeover and giveback of the affected systems. 2. Remove the disk for at least 20 seconds, and then reinsert it. Note: If the disk is part of an aggregate, this option will result in a degraded RAID group and disk reconstruction will be initiated. Otherwise: 1. Verify cabling. 2. Contact NetApp technical support.

## Syslog Message

Multiple disks detected with the same name %s. Disks: %s (%s).

## Parameters

**disk\_name** (STRING): Name of the disks.

**disk\_uid\_list** (STRING): UIDs of disks with the same name.

**disk\_serialno\_list** (STRING): Serial numbers of disks with the same name.

# disk.dynamicqual events

## disk.dynamicqual.fail.parse

### Severity

ERROR

### Description

This message occurs when dynamic device qualification fails due to an invalid `/etc/qual_devices` file. The system is unable to properly recognize disk drives or fully enable all resiliency and performance features.

### Corrective Action

For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "[disk.dynamicqual.fail.parse]" keyword. Request an updated `qual_devices` file from NetApp.

### Syslog Message

Device qualification information file (`/etc/qual_devices`) is invalid. The following error, "%s" has been detected. For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "[disk.dynamicqual.fail.parse]" keyword.

### Parameters

**errorstring** (STRING): Error string for parsing error in `/etc/qual_devices` file.

## disk.dynamicqual.failure.cleared

### Severity

NOTICE

### Description

This message occurs when a previous dynamic device qualification error has been corrected.

### Corrective Action

(None).

## Syslog Message

Device qualification failure condition has been corrected. All devices have been qualified.

## Parameters

(None).

## disk.dynamicqual.failure.invalidFile

### Severity

ALERT

### Description

This message occurs when dynamic device qualification fails due to an invalid /etc/qual\_devices file.

### Corrective Action

See the qual\_devices man page for information about updating device qualification data. Request an updated qual\_devices file from NetApp.

## Syslog Message

Device qualification information file (/etc/qual\_devices) is invalid. See the qual\_devices man page for corrective action. This problem must be corrected within %d hour(s) to avoid a forced system shutdown. The following disk(s) remain unqualified: %s.

## Parameters

**time** (INT): Time until forced shutdown.  
**disklist** (STRING): List of disks that are still not qualified.

## disk.dynamicqual.failure.missingFile

### Severity

ALERT

### Description

This message occurs when dynamic device qualification fails with a missing qualification file.

### Corrective Action

See the qual\_devices man page for information about updating device qualification data. Request a qual\_devices file from NetApp.

## Syslog Message

Device Qualification information file (/etc/qual\_devices) is missing. See the qual\_devices man page for corrective action. This problem must be corrected within %d hour(s) to avoid a forced system shutdown. The following disk(s) remain unqualified: %s.

## Parameters

**time** (INT): Time until forced shutdown.  
**disklist** (STRING): List of disks that are still not qualified.

## disk.dynamicqual.failure.occurred

**Severity**

ALERT

**Description**

This message occurs when dynamic device qualification fails with an unknown error.

**Corrective Action**

See the `qual_devices` man page for information about updating device qualification data.

**Syslog Message**

Device qualification failure occurred. See the `qual_devices` man page for corrective action. This problem must be corrected within %d hour(s) to avoid a forced system shutdown. The following disk(s) remain unqualified: %s.

**Parameters**

**time** (INT): Time until forced shutdown.

**disklist** (STRING): List of disks that are still not qualified.

## **disk.dynamicqual.failure.shutdown**

**Severity**

EMERGENCY

**Description**

This message occurs when one or more disks are not recognized by Data ONTAP®. This can result in a forced system shutdown if the offending disks are not removed or recognized properly within the allowed window of time.

**Corrective Action**

Download the latest Drive Qualification Package (DQP) from NetApp technical support, and then install it on your system.

**Syslog Message**

Device qualification failure has not been corrected for the past %d hour(s). The following disk(s) remain unqualified: %s. System is SHUTTING DOWN.

**Parameters**

**time** (INT): Time until forced shutdown.

**disklist** (STRING): List of disks that are still not qualified.

## **disk.encrypt events**

### **disk.encrypt.destroy.complete**

**Severity**

INFORMATIONAL

**Description**

This message occurs when the disk encrypt destroy operation is completed.

## Corrective Action

(None).

## Syslog Message

Disk %s [S/N %s] has completed encrypt destroy.

## Parameters

**diskName** (STRING): Name of the device.

**serialno** (STRING): Serial number of the device.

## disk.encrypt.revert.complete

### Severity

NOTICE

### Description

This message occurs when the disk encrypt revert operation is completed.

### Corrective Action

(None).

## Syslog Message

Disk %s [S/N %s] has completed encrypt revert.

## Parameters

**diskName** (STRING): Name of the device.

**serialno** (STRING): Serial number of the device.

## disk.encryptaccesserr events

### disk.encryptAccessErr

### Severity

ERROR

### Description

This message occurs when a disk reports a data protection error, and ONTAP® cannot restore permission to access the data through storage encryption methods.

### Corrective Action

If the disk is known to be an encrypting disk, then ONTAP might not have the correct storage encryption passphrase for this device. If possible, import the proper passphrase. If the passphrase cannot be recovered, or if the disk is not an encrypting disk, then all data on the disk is permanently inaccessible. You can return an encrypting disk to service as a spare through the use of the 'disk encrypt sanitize' command. If the device cannot be recovered, then it must be removed from the system.

## Syslog Message

Unable to restore data access on encrypting disk %s; status %s (0x%x).

## Parameters

**diskName** (STRING): Name of the disk.  
**status\_string** (STRING): Status code description.  
**status** (INT): Status code from the SCSI driver.

# disk.encryptaccessfwerr events

## disk.encryptAccessFWErr

### Severity

ERROR

### Description

This message occurs when a disk reports a SCSI data protection error on an attempt to download firmware and ONTAP® cannot restore permission to access the firmware through storage encryption methods.

### Corrective Action

If the disk is known to be an encrypting disk, then ONTAP might not have the correct storage encryption passphrase for this device. If possible, import the proper passphrase. If the passphrase cannot be recovered, or if the disk is not an encrypting disk, then the firmware cannot be downloaded. If the device cannot be recovered, then it must be removed from the system.

### Syslog Message

Unable to restore firmware download access on encrypting disk %s; status %s (0x%x).

## Parameters

**diskName** (STRING): Name of the disk.  
**status\_string** (STRING): Status code description from the SCSI driver.  
**status** (INT): Status code from the SCSI driver.

# disk.encryptcmdfailed events

## disk.encryptCmdFailed

### Severity

ERROR

### Description

This message occurs when the indicated "storage encryption disk" command fails on an encrypting disk that supports storage encryption methods.

### Corrective Action

Verify that storage encryption key servers, if required, are available and functioning, or that expected on-board keys are present.

### Syslog Message

Encrypting disk %s failed disk encrypt %s command with error status %s (0x%x).



## Parameters

**disk** (STRING): Name of the disk.  
**command** (STRING): Storage encryption disk command.  
**error\_string** (STRING): Error status code description.  
**error** (INT): Error status code.

# disk.encryptfipspassphnosync events

## disk.encryptFIPSPassphNoSync

### Severity

NOTICE

### Description

This message occurs when a system with NetApp® Storage Encryption (NSE) enabled detects that the authentication keys (AKs, or passphrases) for the internal FIPS-compliance authorities in a FIPS-certified drive are inconsistent. This can occur when a previous attempt to modify the FIPS-compliance passphrases failed.

### Corrective Action

Use the "security key-manager" commands to display the available AKs from external KMIP key-management servers or onboard key management (OKM). Use the "storage encryption disk modify -fips -key-id" command to assign the desired AK to the FIPS-compliance authorities in the drive. If the condition persists, either replace the drive, or use the "storage encryption disk revert-to-original-state" command to reset the drive to its original condition. Warning: This command also removes all user data from the drive.

### Syslog Message

The FIPS-compliance authorities for NSE drive %s have inconsistent authentication keys (AKs).

### Parameters

**diskName** (STRING): Name of the drive.

# disk.encryptgetmsidfailed events

## disk.encryptGetMSIDFailed

### Severity

NOTICE

### Description

This message occurs when an encrypting disk reports a SCSI data protection error, but Data ONTAP® cannot obtain the Manufacturer's Secure ID (MSID) that is mandatory on disks that support storage encryption methods. Some conditions can cause this error temporarily.

### Corrective Action

If the error persists, replace the disk or remove it from the system.

### Syslog Message

Unable to obtain MSID from disk %s; status "%s" (0x%x).

## Parameters

**diskName** (STRING): Name of the disk.

**status\_string** (STRING): Status code description from the SCSI driver.

**status** (INT): Status code from the SCSI driver.

# disk.encryptnosupport events

## disk.encryptNoSupport

### Severity

ALERT

### Description

This message occurs when ONTAP® software detects one or more FIPS-certified self-encrypting drives, but the environment variable 'bootarg.storageencryption.support' is not present or is set to "false" to indicate that FIPS-certified drives are not supported.

### Corrective Action

Remove all FIPS-certified drives.

### Syslog Message

Detected FIPS-certified encrypting drive %s, but FIPS drives are not supported on this node. %d of %d disks checked are FIPS-certified.

## Parameters

**disk** (STRING): Name of a FIPS-certified drive.

**FIPS\_disks** (INT): Number of FIPS-certified drives detected.

**disk\_count** (INT): Total number of drives checked. For a high-availability pair, this might be fewer than the total number of drives.

# disk.encryptprotdiff events

## disk.encryptProtDiff

### Severity

NOTICE

### Description

This message occurs when self-encrypting drives have different settings of data-at-rest protections; that is, some are protected, having a nondefault Data Key ID, and some are not, having one of the default Data Key IDs. Best protection is provided when all encrypting drives are protected, including spares that might be assigned automatically to an aggregate.

### Corrective Action

Use the "storage encryption disk show -protection-mode !data" command to display the drives that are not protected; use the "storage encryption disk show -protection-mode data" command to show the drives that are protected. Use the "storage encryption disk modify -data-key-id" command to set the Data Key ID to a nondefault value.

## Syslog Message

Self-encrypting drives on the system have dissimilar data protection settings.

## Parameters

(None).

# disk.encryptsupportreqd events

## disk.encryptSupportReqd

### Severity

ALERT

### Description

This message occurs when ONTAP® software detects one or more nonencrypting drives or non-FIPS-certified self-encrypting drives (SEDs), but the environment variable 'bootarg.storageencryption.support' is set to "true" to indicate that all drives must be FIPS-certified.

### Corrective Action

Remove all non-FIPS-certified SEDs and standard drives.

## Syslog Message

Detected nonencrypting drive or non-FIPS-certified drive %s, but only FIPS-certified drives are supported on this node. %d of %d drives checked are FIPS-certified.

## Parameters

**disk** (STRING): Name of a drive that is not FIPS-certified.

**encrypting\_disks** (INT): Number of FIPS-certified drives detected.

**disk\_count** (INT): Total number of drives checked. For a high-availability pair, this might be fewer than the total number of drives.

# disk.eos events

## disk.EOS.OS.error

### Severity

ALERT

### Description

This message occurs when a disk is detected that is not supported in this version of Data ONTAP® or later

### Corrective Action

Remove the disk and replace it with a disk that is supported.

## Syslog Message

%s support ended in version of %s. Current version is %s.

## Parameters

**disk\_information** (STRING): Disk information.

**EOS\_OS\_version** (STRING): Latest version of Data ONTAP supported by the disk.

**cur\_OS\_version** (STRING): Current version of Data ONTAP.

## disk.exceed events

### disk.exceed.cmd.completeTime

#### Severity

INFORMATIONAL

#### Description

This event is generated during disk qualification testing when READ/WRITE I/O requests exceed their expected maximum completion time.

#### Corrective Action

(None).

#### Syslog Message

Disk %s has exceeded %d ms limit — required %llu ms to complete command %s.

#### Parameters

**diskName** (STRING): Name of the disk.

**timeLimit** (INT): Time limit to complete command.

**timeTaken** (LONGINT): Time taken to complete command.

**commandSTR** (STRING): Disk qualification command in string.

## disk.fail events

### disk.fail.ssdstats

#### Severity

INFORMATIONAL

#### Description

This message occurs when a disk is marked failed, is being sanitized, or has entered Maintenance Center.

#### Corrective Action

(None).

#### Syslog Message

Disk %s (%s) failed with rated life used %, percent spare blocks %, spare blocks %s.

#### Parameters

**diskName** (STRING): Name of the disk.

**serialno** (STRING): Serial number.

**ratedLife** (STRING): Rated life used.

**percentSpare** (STRING): Percent spare blocks consumed.

**spareBlks** (STRING): Spare blocks consumed limit.

# disk.failmsg events

## disk.failmsg

### Deprecated

Deprecated as of ONTAP 9.5 - replaced by disk.outOfService.

### Severity

NOTICE

### Description

This message occurs when a drive is marked failed, is being sanitized, or has entered Maintenance Center.

### Corrective Action

(None).

### Syslog Message

Disk %s (%s)%s. %d %s

### Parameters

**diskName** (STRING): Name of the drive.

**serialno** (STRING): Serial number.

**reason** (STRING): Reason for the failure.

**powerOnTime** (INTEX): Time, in minutes, that the drive was powered on when the failure occurred.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

# disk.fakereassignsuccess events

## disk.fakereassignSuccess

### Severity

NOTICE

### Description

This message occurs when a fake reassign is detected on a disk and the sector verify succeeds.

### Corrective Action

(None).

### Syslog Message

disk %s: A fake reassign was issued on sector %llu for disk %s. Disable with boot arg disable-disk-fake-reassign? true.

### Parameters

**diskName** (STRING): Name of the disk.

**sector** (LONGINT): Sector where the error was detected.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

# disk.fli events

## disk.fli.abort.ioSent

### Severity

NOTICE

### Description

This message occurs when an I/O operation below the SCSI layer is aborted because a Foreign LUN Import (FLI) abort message was issued.

### Corrective Action

(None).

### Syslog Message

An I/O with the opcode %s has been aborted on disk %s.

### Parameters

**opcode** (STRING): Opcode that specifies the operation that was aborted.

**diskName** (STRING): Name of the disk.

## disk.fli.abort.ioUnSent

### Severity

NOTICE

### Description

This message occurs when an I/O operation in the disk class layer is aborted because a Foreign LUN Import (FLI) abort message was issued.

### Corrective Action

(None).

### Syslog Message

An I/O with the block count %llu starting at offset %llu has been aborted on disk %s.

### Parameters

**blockCount** (LONGINT): Number of blocks that this I/O operation spans.

**offset** (LONGINT): Start of the I/O operation.

**diskName** (STRING): Name of the disk.

## disk.fli.tmf.complete

### Severity

NOTICE

### Description

This message occurs when a Foreign LUN Import (FLI) message (abort or LUN reset) is completed.

### Corrective Action

(None).

### Syslog Message

Foreign LUN Import (FLI) message: %s completed on disk %s with status %d.

### Parameters

**msgType** (STRING): Specifies which type of FLI message: either abort or LUN reset.

**diskName** (STRING): Name of the disk to which this FLI message was sent.

**returnStatus** (INT): Completion status of the FLI message.

## disk.fw events

### disk.fw.autodownrev

#### Severity

INFORMATIONAL

#### Description

This message occurs when the system detects disks with firmware that is not up-to-date. The background firmware update program will update the firmware automatically if the correct firmware files are in place.

#### Corrective Action

Download the correct firmware file onto the system. Check your support provider's web site or knowledgebase for more information about obtaining the latest disk firmware.

#### Syslog Message

%d disks have downrev firmware. They will be updated automatically using background disk firmware update if the correct firmware files are in place.

#### Parameters

**driveCount** (INT): Number of disks that need to have updated firmware.

### disk.fw.downrevWarning

#### Severity

ERROR

#### Description

This message occurs when the system detects disks with firmware that is not up-to-date.

#### Corrective Action

Download and install the latest disk firmware file. Check your support provider's web site or knowledgebase for more information about obtaining the latest disk firmware.

#### Syslog Message

%d disks have downrev firmware that you need to update.

#### Parameters

**driveCount** (INT): Number of disks that need to have updated firmware.

# disk.healthtrigger events

## disk.healthTrigger

### Severity

NOTICE

### Description

This message occurs when the disk returns a health trigger event.

### Corrective Action

No action is required because this message is informational only. If the disk is having problems processing I/O, an error will be returned.

### Syslog Message

Disk %s received NHT health trigger (0x%x 0x%x 0x%x 0x%x)

### Parameters

**diskName** (STRING): Name of the device.  
**sense\_key** (INTHEX): SCSI sense key.  
**sense\_code** (INTHEX): SCSI sense code.  
**qualifier** (INTHEX): SCSI sense code qualifier.  
**fru\_failed** (INTHEX): SCSI FRU code.

# disk.init events

## disk.init.badSectorSize

### Severity

ERROR

### Description

This message occurs when a disk reports a sector size that does not match the drive qualification table.

### Corrective Action

Replace the disk.

### Syslog Message

Disk %s has an unexpected sector size (%d bytes) and cannot be used.

### Parameters

**diskName** (STRING): Name of the disk.  
**sector\_size** (INT): Invalid sector size.

## disk.init.err.capacityFlash

### Severity

ERROR



## Description

This message occurs when a capacity-optimized flash solid-state drive (SSD) is detected by a node that does not support this type of SSD.

## Corrective Action

Remove the disk from the system.

## Syslog Message

Disk %s cannot be used on a node that does not support capacity-optimized flash.

## Parameters

**diskName** (STRING): Name of the disk.

## disk.init.err.flashOptimized

### Severity

ERROR

### Description

This message occurs when a hard disk drive (HDD) is detected by a node configured with the All-Flash Optimized personality. With this personality, only solid-state drives (SSD) can be used.

### Corrective Action

Remove the disk from the system.

### Syslog Message

Disk %s cannot be used on a node configured as All-Flash Optimized.

### Parameters

**diskName** (STRING): Name of the disk.

## disk.init.err.unsupp.shelf

### Severity

ERROR

### Description

This message occurs when an unsupported NVMe solid-state drive (SSD) is discovered in an external NVMe shelf. This NVMe SSD is only supported as an internal drive on the A800 system.

### Corrective Action

Remove the drive from the external NVMe shelf.

### Syslog Message

Drive %s cannot be used in an external NVMe shelf.

### Parameters

**disk\_name** (STRING): Name of the SSD.

**disk\_information** (STRING): Formatted information about the SSD. This includes the drive's vendor, model, firmware revision, and serial number.

## **disk.init.err.unsupported**

### **Severity**

ERROR

### **Description**

This message occurs when an unsupported disk is detected on a node.

### **Corrective Action**

Remove the disk from the system.

### **Syslog Message**

Disk %s is not recognized as a supported disk and cannot be used on this node.

### **Parameters**

**diskName** (STRING): Name of the disk.

## **disk.init.err.zns**

### **Severity**

ERROR

### **Description**

This message occurs when a zoned namespace solid-state drive (SSD-ZNS) is detected by a node that does not support this type of SSD.

### **Corrective Action**

Remove the disk from the system.

### **Syslog Message**

Disk "%s" cannot be used on a node that does not support SSD-ZNS disks.

### **Parameters**

**diskName** (STRING): Name of the disk.

## **disk.init.error.capacity**

### **Severity**

ERROR

### **Description**

This message occurs when a disk reports an unexpected capacity.

### **Corrective Action**

Replace the disk.

### **Syslog Message**

Disk %s has an unexpected capacity (%llu sectors) and cannot be used. Replace it.

### **Parameters**

**diskName** (STRING): Name of the disk.

**lastSector** (LONGINT): Last sector of the disk.  
**productID** (STRING): Product name of the disk.

## **disk.init.error.serialno**

### **Severity**

ERROR

### **Description**

This message occurs when the inquiry command does not return a valid serial number or device ID for the device. The device is failed during initialization.

### **Corrective Action**

Replace the disk.

### **Syslog Message**

Serial number or device ID not available for %s.

### **Parameters**

**disk\_name** (STRING): Disk name.  
**page\_code** (INT): Page code that was sent.  
**error\_pcode** (INT): Page code received in error.

## **disk.init.failure.error**

### **Severity**

ERROR

### **Description**

This message occurs when a drive fails initialization. For NVMe drives, the returned error information is translated into the SCSI equivalent.

### **Corrective Action**

Replace the drive.

### **Syslog Message**

Drive %s failed initialization due to error %d, sense code(%x %x %x %x).

### **Parameters**

**diskName** (STRING): Name of the drive.  
**errorCode** (INT): Internal E\_SCSI return code.  
**sense\_key** (INTHEX): Sense key.  
**sense\_code** (INTHEX): Additional sense code.  
**qualifier** (INTHEX): Additional sense code qualifier.  
**fru\_failed** (INTHEX): FRU code.

## **disk.init.failure.spinup**

### **Severity**

ERROR

**Description**

This message occurs when a disk cannot spin up during initialization. The disk is marked as failed and is not used.

**Corrective Action**

Removed the disk from the system.

**Syslog Message**

Disk %s has failed to spin up and cannot be used. Replace it with a new drive.

**Parameters**

**diskName** (STRING): Name of the disk.

**disk.init.failureBytes****Severity**

ERROR

**Description**

This message occurs when the system discovers a previously failed disk. The system will not use the disk and it should be replaced.

**Corrective Action**

Replace the failed disk.

**Syslog Message**

Failed disk %s detected during disk initialization.

**Parameters**

**diskName** (STRING): Name of the disk.

**disk.init.invalidDDR****Severity**

INFORMATIONAL

**Description**

This message occurs when a disk contains a Dynamic Drive Recognition (DDR) label entry that is not for this disk. The system rewrites the DDR label with proper values.

**Corrective Action**

(None).

**Syslog Message**

Disk %s has an invalid DDR entry. DDR label will be corrected and rewritten automatically.

**Parameters**

**diskName** (STRING): Name of the disk.

**vendorID** (STRING): Vendor name of the disk.

**productID** (STRING): Product name of the disk.

## disk.init.protype.incompat

### Severity

ERROR

### Description

This message occurs when a hard disk drive (HDD) has a protection type enabled that is not supported. Protection type is a disk option that determines how information is physically saved on the disk.

### Corrective Action

Remove the disk from the system.

### Syslog Message

Disk %s has protection type %x enabled and cannot be used. %s

### Parameters

**diskName** (STRING): Name of the disk.

**protType** (INHEX): Value indicating the version of the enabled protection type.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.init.recognizedDDR

### Severity

INFORMATIONAL

### Description

This message occurs when a Dynamic Drive Recognition (DDR) label on a drive with dynamically qualified data is properly updated with aliased data.

### Corrective Action

(None).

### Syslog Message

Disk %s has been updated with the proper drive name alias. The DDR label will be rewritten automatically.

### Parameters

**diskName** (STRING): Name of the disk.

**vendorID** (STRING): Vendor name of the disk.

**productID** (STRING): Product name of the disk.

## disk.init.unknownDisk

### Severity

ERROR

### Description

This message occurs when a unique identifier cannot be created for a disk due to internal disk problems.

### Corrective Action

The disk needs to be replaced, remove it from the system.

## Syslog Message

Unable to identify disk %s.

## Parameters

**diskName** (STRING): Name of the disk.

# disk.iomediumerror events

## disk.ioMediumError

### Severity

NOTICE

### Description

This message occurs when the drive reports a medium error.

### Corrective Action

(None).

## Syslog Message

Medium error on disk %s: op %s sector %llu %s - %s %s (%x %x %x %x) (%d) %s

## Parameters

**diskName** (STRING): Name of the disk.

**op** (STRING): I/O operation being performed.

**sector** (LONGINT): Sector where the error was detected.

**senseInfo** (STRING): Sense data.

**sCode** (STRING): Sector code.

**disk\_info** (STRING): Static description of events on the disk.

**sense\_key** (INTHEX): Actual sense key.

**sense\_code** (INTHEX): Actual sense code.

**qualifier** (INTHEX): Actual sense code qualifier.

**fru\_failed** (INTHEX): Actual FRU code.

**CTime** (INT): Time, in milliseconds, from when the command was first issued until this I/O operation completed.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

# disk.ioReassignfailed events

## disk.ioReassignFailed

### Severity

ALERT

### Description

This message occurs when a disk reports a medium error and the sector reassignment fails. The system will fail the disk after possible recovery of retrievable disk data.

### Corrective Action

Replace the disk when the system indicates that you should remove it in subsequent logged events.

### Syslog Message

disk %s: sector %llu failed to be reassigned (%d). %s

### Parameters

**diskName** (STRING): Name of the disk.

**sector** (LONGINT): Sector where the error was detected.

**ETime** (INT): Time, in milliseconds, from when the command was first issued until this I/O operation failed.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.ioReassignsuccess events

### disk.ioReassignSuccess

#### Severity

NOTICE

#### Description

This message occurs when a medium error is detected on a disk and the sector reassignment succeeds.

#### Corrective Action

(None).

#### Syslog Message

disk %s: sector %llu was reassigned (%d). %s

#### Parameters

**diskName** (STRING): Name of the disk.

**sector** (LONGINT): Sector where the error was detected.

**CTime** (INT): Time, in milliseconds, from when the command was first issued until this I/O operation completed.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.ioRecoverederror events

### disk.ioRecoveredError.pfa

#### Severity

ERROR

#### Description

This message occurs when a disk determines that it will fail shortly. This occurs when a threshold internal to the disk indicates that a failure is imminent.

## Corrective Action

To avoid an uncontrolled failure, replace the disk as soon as possible. If the disk is a member of a RAID volume, the disk will be failed by RAID when this error occurs on a RAID I/O and the RAID group is not in degraded mode.

## Syslog Message

Recovered error predictive failure alert on disk %s: op %s sector %llu %s - %s (%x %x %x %x)

## Parameters

**diskName** (STRING): Name of the disk.

**op** (STRING): I/O operation being performed.

**sector** (LONGINT): Sector where the error was detected.

**senseInfo** (STRING): Sense data.

**disk\_info** (STRING): Static description of events on the disk.

**sense\_key** (INTHEX): Actual sense key.

**sense\_code** (INTHEX): Actual sense code.

**qualifier** (INTHEX): Actual sense code qualifier.

**fru\_failed** (INTHEX): Actual FRU code.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.ioRecoveredError.reassign

### Severity

NOTICE

### Description

This message occurs when a disk automatically reassigns a sector on a disk.

### Corrective Action

(None).

### Syslog Message

Recovered error on disk %s: op %s sector %llu %s - %s (%x %x %x %x) %s

### Parameters

**diskName** (STRING): Name of the disk.

**op** (STRING): I/O operation being performed.

**sector** (LONGINT): Sector where the error was detected.

**senseInfo** (STRING): Sense data.

**disk\_info** (STRING): Static description of events on the disk.

**sense\_key** (INTHEX): Actual sense key.

**sense\_code** (INTHEX): Actual sense code.

**qualifier** (INTHEX): Actual sense code qualifier.

**fru\_failed** (INTHEX): Actual FRU code.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.ioRecoveredError.retry



**Severity**

INFORMATIONAL

**Description**

This message occurs when a recovered error is detected on a disk.

**Corrective Action**

(None).

**Syslog Message**

Recovered error on disk %s: op %s sector %llu %s - %s (%x %x %x %x) (%d) %s

**Parameters**

**diskName** (STRING): Name of the disk.

**op** (STRING): I/O operation being performed.

**sector** (LONGINT): Sector where the error was detected.

**senseInfo** (STRING): Sense data.

**disk\_info** (STRING): Static description of events on the disk.

**sense\_key** (INTHEX): Actual sense key.

**sense\_code** (INTHEX): Actual sense code.

**qualifier** (INTHEX): Actual sense code qualifier.

**fru\_failed** (INTHEX): Actual FRU code.

**ETime** (INT): Time, in milliseconds, from when the command was first issued until this I/O operation failed.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

## disk.lipstormdetect events

### disk.lipStormDetect

**Severity**

NOTICE

**Description**

This message occurs when a Fibre Channel disk detects a large number of Loop Initialization Procedure requests (a LIP storm) and has reset itself.

**Corrective Action**

None required. If this happens a second time, the disk will be failed.

**Syslog Message**

shm: disk %s reports a LIP storm reset. %s. Error %x/%x/%x/%x.

**Parameters**

**diskName** (STRING): Name of the disk (host adapter, loop ID).

**disk\_info** (STRING): Static description of events on the disk.

**sense\_key** (INT): Actual sense key.

**sense\_code** (INT): Actual sense code.

**qualifier** (INT): Actual sense code qualifier.

**fru\_failed** (INT): Actual FRU code.

# disk.lostwritedetected events

## disk.lostwriteDetected

### Severity

ERROR

### Description

This message occurs when the system detects a lost write.

### Corrective Action

(None).

### Syslog Message

Disk level lost write detected on WRITE VERIFY: Disk %s, Block #%%llu: Volume %s, Fileid %d, Block #%%llu: Expected 0x%x, Recomputed as 0x%x.

### Parameters

**diskName** (STRING): Name of the device.

**bno** (LONGINT): Block ID.

**vol** (STRING): Volume name.

**fileid** (INT): File ID.

**block** (LONGINT): File block ID.

**cksum** (INTHEX): Expected checksum.

**cksum2** (INTHEX): Computed checksum.

# disk.lun events

## disk.lun.allMemoryAvailable

### Severity

NOTICE

### Description

This message occurs to confirm that enough memory is present to successfully allocate command blocks for all of the disks in the system and for additional LUNS.

### Corrective Action

(None).

### Syslog Message

disk\_init :Allocated all memory required for array LUNs.

### Parameters

(None).

## disk.lun.outOfMemory

### Severity

NOTICE

## Description

This message occurs when the system cannot allocate the additional memory for LUNs.

## Corrective Action

(None).

## Syslog Message

disk\_init: Unable to allocate memory for array LUN commands on %s.

## Parameters

**diskName** (STRING): Name of the device that could not allocate memory.

## disk.lun.sizeTooLarge

### Severity

ALERT

## Description

This message occurs when an array LUN on an external RAID storage subsystem or a virtual disk is too large to be used as storage by the controller.

## Corrective Action

Reduce this array LUN size, or re-create the virtual disk within the parameter limit.

## Syslog Message

Array LUN or virtual disk %s is too large (%s bytes) and cannot be used. The maximum permissible array LUN or virtual disk size is %s. Reduce this array LUN size by %lld bytes, or re-create the virtual disk within the parameter limit (%lld %d byte blocks).

## Parameters

**diskName** (STRING): Name of the array LUN or virtual disk that is too large.

**lunSize** (STRING): Size, expressed in bytes, of the array LUN or virtual disk that is too large.

**maxSize** (STRING): Maximum permissible array LUN or virtual disk size, expressed in bytes.

**byteDiff** (LONGINT): Difference, expressed in bytes, between the array LUN or virtual disk size and the maximum size Data ONTAP® supports.

**blockDiff** (LONGINT): Difference, expressed in blocks, between the array LUN or virtual disk size and the maximum size Data ONTAP supports.

**blockSize** (INT): Block size of the array LUN or virtual disk.

## disk.lun.sizeTooSmall

### Severity

ALERT

## Description

This message occurs when an array LUN on an external RAID storage subsystem or a virtual disk is too small to be used as storage by the controller.

## Corrective Action

Delete the array LUN or virtual disk and re-create it within the parameter limit.

## Syslog Message

LUN or virtual disk %s is too small (%s bytes) and cannot be used. The minimum permissible LUN or virtual disk size is %s. Increase this LUN or virtual disk size by %lld bytes (%lld %d byte blocks).

## Parameters

**diskName** (STRING): Name of the LUN or virtual disk that is too small.

**lunSize** (STRING): Size, in bytes, of the LUN or virtual disk that is too small.

**minSize** (STRING): Minimum permissible LUN or virtual disk size, expressed in bytes.

**byteDiff** (LONGINT): Difference, in bytes, between the LUN or virtual disk size and the minimum size Data ONTAP® supports.

**blockDiff** (LONGINT): Difference, in blocks, between the LUN or virtual disk size and the minimum size Data ONTAP supports.

**blockSize** (INT): Block size of the LUN or virtual disk.

## disk.lun.unsupportedBlkSize

### Severity

ALERT

### Description

This message occurs when an array logical unit (LUN) on the external RAID storage subsystem reports an unsupported block size.

### Corrective Action

Reformat the array LUN with a supported block size.

## Syslog Message

disk\_init: Array LUN %s reports a block size of %u bytes and it cannot be used. Reformat the array LUN with a supported block size.

## Parameters

**diskName** (STRING): Name of the array LUN with the unsupported block size.

**sector\_size** (INT): Unsupported block size.

## disk.maint events

### disk.maint.test.end

### Severity

NOTICE

### Description

This message occurs when maintenance testing has completed.

### Corrective Action

(None).

## Syslog Message

Disk %s Model [%s] S/N [%s] has completed maintenance testing with status %d (0x%x,0x%x,0x%x,0x%x).

## Parameters

**disk\_name** (STRING): Name of the disk.

**productID** (STRING): Product name of the disk.

**serialno** (STRING): Serial number of the disk.

**status** (INT): Maintenance test completion status where 0 represents successful completion of the test.

**fa0** (INTHEX): Failure byte 0.

**fa1** (INTHEX): Failure byte 1.

**fa2** (INTHEX): Failure byte 2.

**fa3** (INTHEX): Failure byte 3.

## disk.maint.test.start

### Severity

NOTICE

### Description

This message occurs when maintenance testing begins on a disk.

### Corrective Action

(None).

### Syslog Message

Disk %s Model [%s] S/N [%s] has started maintenance testing.

## Parameters

**disk\_name** (STRING): Name of the disk.

**productID** (STRING): Product name of the disk.

**serialno** (STRING): Serial number of the disk.

## disk.max events

## disk.max.partitions

### Severity

ERROR

### Description

This message occurs when ONTAP® software cannot partition a disk because the system already has the maximum number of partitioned disks.

### Corrective Action

Contact Contact NetApp technical support. for assistance.

### Syslog Message

Disk %s cannot be partitioned as the system already has %d partitioned disks.

## Parameters

**diskName** (STRING): Name of the disk.

**maxPartitionLimit** (INT): Maximum number of disks which can be partitioned.

# disk.min events

## disk.min.OS.error

### Severity

ERROR

### Description

This message occurs when a disk is detected that is not supported in this version of Data ONTAP®.

### Corrective Action

Remove the disk and replace it with a disk that is supported, or upgrade the kernel to the required version.

### Syslog Message

%s is not supported because it requires a minimum kernel version of %s. Current version is %s.

### Parameters

**disk\_information** (STRING): Disk information.

**min\_OS\_version** (STRING): Disk's minimum supported Data ONTAP version.

**cur\_OS\_version** (STRING): Current version of the Data ONTAP.

# disk.outofservice events

## disk.outOfService

### Severity

NOTICE

### Description

This message occurs when a drive is removed from service because it has been marked failed, is being sanitized, or has entered Maintenance Center.

### Corrective Action

(None).

### Syslog Message

Drive %s (%s)%s. Power-On Hours: %s, GList Count: %d, Drive Info: %s.

### Parameters

**diskName** (STRING): Name of the drive.

**serialno** (STRING): Serial number of the drive.

**reason** (STRING): Reason for removal from service.

**powerOnHours** (STRING): Time, in hours, that the drive was powered on when the failure occurred.

**glistEntries** (INT): Number of grown defect list entries.

**disk\_information** (STRING): Formatted information about the drive. This includes the drive's vendor, model, firmware revision, and serial number.

# disk.partdisknotsuppmcc events

## disk.partDiskNotSuppMCC

### Severity

NOTICE

### Description

This message occurs when a partitioned disk is found in a MetroCluster™ configuration.

### Corrective Action

Remove the partitioned disk, because it is not supported in a MetroCluster configuration.

### Syslog Message

Ignoring disk %s because partitioned disks are not supported in a MetroCluster™ configuration.

### Parameters

**diskName** (STRING): Name of the disk.

## disk.partition events

### disk.partition.exceeded

### Severity

ERROR

### Description

This message occurs when the disk partition layout contains more partitions than are supported by the system.

### Corrective Action

Fix the unsupported disk partition layout by using the 'disk unpartition' command to make the disk a spare disk.

### Syslog Message

Disk %s has %d partitions, which exceeds the maximum of %d partitions.

### Parameters

**disk\_name** (STRING): Name of the disk.

**partition\_count** (INT): Number of partitions on the disk.

**max\_partition\_count** (INT): Maximum number of partitions supported per disk.

## disk.partner events

### disk.partner.encrRevertDone

### Severity

NOTICE

### Description

This message occurs when a node receives notice that a partner node has completed a disk encrypt revert operation.

**Corrective Action**

(None).

**Syslog Message**

The partner has completed reverting drive %s to its original state.

**Parameters**

**diskName** (STRING): Name of the device.

**disk.partner.encrRevertStart****Severity**

NOTICE

**Description**

This message occurs when a node receives notice that a partner node has started a disk encrypt revert operation.

**Corrective Action**

(None).

**Syslog Message**

The partner has started reverting drive %s to its original state.

**Parameters**

**diskName** (STRING): Name of the device.

**disk.partner.sanitizeStart****Severity**

INFORMATIONAL

**Description**

This message occurs when a message is received from the partner for a disk sanitize start.

**Corrective Action**

(None).

**Syslog Message**

The partner has started a disk sanitize on %s.

**Parameters**

**diskName** (STRING): Name of the disk.

**disk.readreservationfailed events****disk.readReservationFailed****Severity**

ERROR



## Description

This message occurs when a SCSI (Small Computer System Interface) read reservation fails.

## Corrective Action

This is not a problem if seen occasionally. If this message is seen repeatedly, or is blocking access to a LUN for longer than 10 minutes, open a support case to investigate the underlying cause.

## Syslog Message

Disk read reservation failed on %s CDB %s - %s (%x %x %x)

## Parameters

**diskName** (STRING): Name of the device.

**cdb** (STRING): CDB (Command Descriptor Block) containnig the I/O operation being performed.

**sSenseKey** (STRING): Sense key value.

**iSenseKey** (INTHEX): Actual sense key.

**iASC** (INTHEX): SCSI Additional Sense Code.

**iASCQ** (INTHEX): SCSI Additional Sense Code Qualifier.

**sm** (STRING): SCSI Persistent Reservation state machine.

**state** (INT): SCSI Persistent Reservation state machine step.

**out\_status\_flags** (INTHEX): Status flag contains bits describing which of the following status values are valid.

**ha\_status** (INTHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**target\_status** (INTHEX): Value returned by the target describing completion status of I/O request.

**target\_rsp\_code** (INTHEX): Target response code.

# disk.releasefailed events

## disk.releaseFailed

### Severity

ERROR

### Description

This message occurs when a SCSI release reservation fails.

### Corrective Action

This is not a problem if seen occasionally. If this message is seen repeatedly, or is blocking access to a LUN for longer than 10 minutes, open a support case to investigate the underlying cause.

### Syslog Message

Disk release failed on %s CDB %s - %s (%x %x %x)

### Parameters

**diskName** (STRING): Name of the device.

**cdb** (STRING): CDB (Command Descriptor Block) containing the I/O operation being performed.

**sSenseKey** (STRING): Sense key value.

**iSenseKey** (INTHEX): Actual sense key.

**iASC** (INTHEX): SCSI Additional Sense Code.

**iASCQ** (INTHEX): SCSI Additional Sense Code Qualifier.

**sm** (STRING): SCSI Persistent Reservation state machine.

**state** (INT): SCSI Persistent Reservation state machine step.

**out\_status\_flags** (INTEX): Status flag contains bits describing which of the following status values are valid.

**ha\_status** (INTEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**target\_status** (INTEX): Value returned by the target describing completion status of I/O request.

**target\_rsp\_code** (INTEX): Target response code.

## disk.reservedelay events

### disk.reserveDelay

#### Severity

NOTICE

#### Description

This message occurs when a SCSI reservation command has a longer latency than expected during takeover. This delay could cause a system outage in which case the disk should be removed to prevent further problems. .

#### Corrective Action

If you see this message before an outage on takeover, remove the specified disk.

#### Syslog Message

The disk reservation command sent to %s has taken longer than expected: %d msec during takeover.

#### Parameters

**diskName** (STRING): Name of the device.

**elapsed\_time** (INT): Number of milliseconds for the command to complete.

## disk.reserveerror events

### disk.reserveError

#### Severity

NOTICE

#### Description

This message occurs when a SCSI reservation is requested for a disk or LUN that is no longer seen.

#### Corrective Action

A disk or LUN might have been removed from the storage subsystem or the path to a disk or LUN has been lost. Verify that all paths to the disk or LUN are visible to the appliance using the storage show disk -p command.

#### Syslog Message

disk\_reserve: %s was not found: CDB %s - %s (%x %x %x)

#### Parameters

**diskName** (STRING): Name of the device.

**cdb** (STRING): CDB (Command Descriptor Block) containing the I/O operation being performed.

**sSenseKey** (STRING): Sense key string.

**iSenseKey** (INTHEX): Sense key value.

**iASC** (INTHEX): SCSI Additional Sense Code.

**iASCQ** (INTHEX): SCSI Additional Sense Code Qualifier.

**sm** (STRING): SCSI Persistent Reservation state machine.

**state** (INT): SCSI Persistent Reservation state machine step.

**out\_status\_flags** (INTHEX): Status flag containing bits describing which of the following status values are valid.

**ha\_status** (INTHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**target\_status** (INTHEX): Value returned by the target describing completion status of I/O request.

**target\_rsp\_code** (INTHEX): Target response code.

## disk.reservefailed events

### disk.reserveFailed

#### Severity

ERROR

#### Description

This message occurs when a SCSI reservation fails.

#### Corrective Action

This is not a problem if seen occasionally. If this message is seen repeatedly, or is blocking access to a LUN for longer than 10 minutes, contact NetApp technical support to open a support case to investigate the underlying cause.

#### Syslog Message

Disk reservation failed on %s CDB %s - %s (%x %x %x)

#### Parameters

**diskName** (STRING): Name of the device.

**cdb** (STRING): CDB (Command Descriptor Block) containing the I/O operation being performed.

**sSenseKey** (STRING): Sense key string.

**iSenseKey** (INTHEX): Actual sense key.

**iASC** (INTHEX): SCSI Additional Sense Code.

**iASCQ** (INTHEX): SCSI Additional Sense Code Qualifier.

**sm** (STRING): SCSI Persistent Reservation state machine.

**state** (INT): SCSI Persistent Reservation state machine step.

**out\_status\_flags** (INTHEX): Status flag containing bits describing which of the following status values are valid.

**ha\_status** (INTHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**target\_status** (INTHEX): Value returned by the target describing completion status of I/O request.

**target\_rsp\_code** (INTHEX): This is the target response code

## disk.rewritedatafailed events

### disk.rewriteDataFailed

**Severity**

ALERT

**Description**

This message occurs when the system encounters an error and cannot rewrite data that was difficult to read on ATA drives.

**Corrective Action**

(None).

**Syslog Message**

Unable to rewrite %s on disk %s.

**Parameters**

**command** (STRING): SCSI command that reported a problem.

**diskName** (STRING): Name of the device.

## disk.sanit events

### disk.sanit.complete

**Severity**

INFORMATIONAL

**Description**

This message occurs when disk sanitization has been completed.

**Corrective Action**

(None).

**Syslog Message**

Disk %s [S/N %s] has completed sanitization.

**Parameters**

**diskName** (STRING): Name of the device.

**serialno** (STRING): Serial number of the device.

### disk.sanit.error

**Severity**

ERROR

**Description**

This message occurs when a disk returns an error during sanitization.

**Corrective Action**

Verify that the drive is capable of sanitization. See the ONTAP® Physical Storage Management Guide <https://library.netapp.com/ecmdocs/ECMP12458210/html/GUID-1DC1DEC2-AAB3-483E-805D-82E722023AA6.html> for more information.

## Syslog Message

Disk %s has received an error %d (0%hx/0%hx/0%hx/0%hx) during sanitization.

### Parameters

**diskName** (STRING): Name of the disk.  
**scsi\_error** (INT): SCSI error code.  
**sense\_key** (INT): Actual sense key.  
**sense\_code** (INT): Actual sense code.  
**qualifier** (INT): Actual sense code qualifier.  
**fru\_failed** (INT): Actual FRU code.

## disk.sanit.formatComplete

### Severity

INFORMATIONAL

### Description

This message occurs when the disk sanitization format phase is complete.

### Corrective Action

(None).

## Syslog Message

Disk sanitization format complete for disk %s [S/N %s], begin pattern writes.

### Parameters

**diskName** (STRING): Name of the device.  
**serialno** (STRING): Serial number of the device.

## disk.sanit.formatError

### Severity

ERROR

### Description

This message occurs when disk sanitization formatting fails.

### Corrective Action

(None).

## Syslog Message

Disk %s has received an error %d (0%hx/0%hx/0%hx) during the formatting of the drive. Format will be retried within fifteen minutes.

### Parameters

**diskName** (STRING): Name of the device.  
**status** (INT): Error code.  
**sense\_key** (INTHEX): Actual sense key.  
**sense\_code** (INTHEX): Actual sense code.  
**qualifier** (INTHEX): Actual sense code qualifier.

## disk.sanit.formatFailed

### Severity

ERROR

### Description

This message occurs when disk sanitization reformatting has been retried unsuccessfully the maximum number of times.

### Corrective Action

Disk sanitization reformatting attempts have been unsuccessful. To make the disk usable, it must be reformatted. You can do this in maintenance mode with the SCSI "format" command.

### Syslog Message

Disk sanitization reformat has failed for disk %s.

### Parameters

**diskName** (STRING): Name of the device.

## disk.sanit.patternDone

### Severity

INFORMATIONAL

### Description

This message occurs when disk sanitization has completed a cycle of a pattern write.

### Corrective Action

(None).

### Syslog Message

Disk %s [S/N %s] cycle %d pattern write of 0x%02x completed in %02d:%02d:%02d.

### Parameters

**diskName** (STRING): Name of the device.

**serialno** (STRING): Serial number of the device.

**cycle** (INT): Sanitization cycle.

**pattern** (INT): Pattern byte.

**hours** (INT): Time in hours.

**minutes** (INT): Time in minutes.

**seconds** (INT): Time in seconds.

## disk.sanit.randomPatternDone

### Severity

INFORMATIONAL

### Description

This message occurs when disk sanitization has completed a cycle of random data write.

## Corrective Action

(None).

## Syslog Message

Disk %s [S/N %s] cycle %d pattern write of random data completed in %02d:%02d:%02d.

## Parameters

**diskName** (STRING): Name of the device.  
**serialno** (STRING): Serial number of the device.  
**cycle** (INT): Sanitization cycle.  
**hours** (INT): Time in hours.  
**minutes** (INT): Time in minutes.  
**seconds** (INT): Time in seconds.

## disk.sanit.reformatComplete

### Severity

INFORMATIONAL

### Description

This message occurs when disk sanitization reformatting is complete.

### Corrective Action

(None).

### Syslog Message

Disk %s has been successfully reformatted. To sanitize disk, restart the disk sanitize command. To bypass format, use -s option.

### Parameters

**diskName** (STRING): Name of the device.

## disk.sanit.sanitComplete

### Severity

INFORMATIONAL

### Description

This message occurs when the SCSI sanitize command phase of a 'disk sanitize' command is complete, and the data pattern write phase begins.

### Corrective Action

(None).

### Syslog Message

Sanitization phase of the 'disk sanitize' command was completed for disk %s [S/N %s]. Pattern writes have begun.

### Parameters

**diskName** (STRING): Name of the device.  
**serialno** (STRING): Serial number of the device.

# disk.selectiontimeout events

## disk.selectionTimeout

### Severity

ERROR

### Description

This message occurs when a selection timeout is received.

### Corrective Action

(None).

### Syslog Message

Disk %s: selection timeout

### Parameters

**diskName** (STRING): Name of the device.

# disk.senseerror events

## disk.senseError

### Severity

ERROR

### Description

This message occurs when the drive reports an error.

### Corrective Action

Check the event log for related error messages and corrective actions.

### Syslog Message

Disk %s: op %s sector %llu %s - %s (%X %X %X %X)

### Parameters

**diskName** (STRING): Name of the disk.

**op** (STRING): I/O operation being performed.

**sector** (LONGINT): Sector where the error was detected.

**senseInfo** (STRING): Sense data.

**sCode** (STRING): Sector code.

**sense\_key** (INTHEX): SCSI sense key.

**sense\_code** (INTHEX): SCSI sense code.

**qualifier** (INTHEX): SCSI sense code qualifier.

**fru\_failed** (INTHEX): SCSI FRU code.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.



# disk.senseinfoinvalid events

## disk.senseInfoNotValid

### Severity

ERROR

### Description

This message occurs when invalid sense information is returned.

### Corrective Action

(None).

### Syslog Message

Disk %s: sense data information field is not valid.

### Parameters

**diskName** (STRING): Name of the device.

# disk.setarraylqd events

## disk.setArrayLQD

### Severity

INFORMATIONAL

### Description

This message occurs when you set the LUN queue depth for all the LUNs of an array.

### Corrective Action

(None).

### Syslog Message

All the LUNs on array '%s' have a LUN queue depth value set to %d.

### Parameters

**arrayName** (STRING): Name of the array.

**lqd** (INT): LUN queue depth value.

# disk.sizegrew events

## disk.sizeGrew

### Severity

NOTICE

### Description

This message occurs when the controller detects an increase in the capacity of a virtual disk.

## Corrective Action

(None).

## Syslog Message

Size of virtual disk %s has been changed from %s to %s, an increase of %lld bytes (%lld %d-byte blocks).

## Parameters

**diskName** (STRING): Name of the virtual disk that has an increased capacity.

**Size** (STRING): Original size of the virtual disk in bytes.

**NewSize** (STRING): Current size of the virtual disk in bytes.

**byteDiff** (LONGINT): Difference in bytes between the original size and the current size of the virtual disk.

**blockDiff** (LONGINT): Difference in blocks between the original size and the current size of the virtual disk.

**blockSize** (INT): Block size in bytes of the virtual disk.

# disk.sizeshrank events

## disk.sizeShrank

### Severity

NOTICE

### Description

This message occurs when the controller detects a decrease in the capacity of a virtual disk.

### Corrective Action

(None).

### Syslog Message

Size of virtual disk %s has changed from %s to %s, a decrease of %lld bytes (%lld %d-byte blocks).

### Parameters

**diskName** (STRING): Name of the virtual disk that has a decreased capacity.

**Size** (STRING): Original size of the virtual disk in bytes.

**NewSize** (STRING): Current size of the virtual disk in bytes.

**byteDiff** (LONGINT): Difference in bytes between the original size and the current size of the virtual disk.

**blockDiff** (LONGINT): Difference in blocks between the original size and the current size of the virtual disk.

**blockSize** (INT): Block size in bytes of the virtual disk.

# disk.spindleerror events

## disk.spindleError.detect

### Severity

ERROR

### Description

This message occurs when a disk returns a spindle detect error. The disk will be recommended for failure.

## Corrective Action

Disk firmware has detected spindle motor problems. The disk will be recommended for failure. If the disk is a member of a RAID volume, it will be failed by RAID when this error occurs on a RAID I/O and the raid group is not in degraded mode.

## Syslog Message

Disk %s has reported a spindle detect error: (%x %x %x %x)

## Parameters

**diskName** (STRING): Name of the disk.  
**sense\_key** (INTHEX): Actual sense key  
**sense\_code** (INTHEX): Actual sense code.  
**qualifier** (INTHEX): Actual sense code qualifier.  
**fru\_failed** (INTHEX): Actual FRU code.

## disk.spindleError.notice

### Severity

INFORMATIONAL

### Description

This message occurs when a disk returns a spindle detect error.

### Corrective Action

Disk firmware has detected spindle motor problems. The option to recommend the drive for failure is not set, so the disk will not be automatically recommended for failure.

### Syslog Message

(None).

### Parameters

**diskName** (STRING): Name of the disk.  
**sense\_key** (INT): Actual sense key.  
**sense\_code** (INT): Actual sense code.  
**qualifier** (INT): Actual sense code qualifier.  
**fru\_failed** (INT): Actual FRU code.

## disk.stalledtoolong events

### disk.stalledTooLong

#### Severity

ERROR

#### Description

This message occurs when a disk has been stalled too long.

#### Corrective Action

(None).

## Syslog Message

I/O for disk %s has been suspended %d times for %d seconds.

## Parameters

**diskName** (STRING): Name of the device.

**count** (INT): Count of times suspended.

**error** (INT): Time suspended.

# disk.timeout.events

## disk.timeout.flush.end

### Severity

NOTICE

### Description

This message occurs when the disk driver stops flushing aggressive timeout I/Os.

### Corrective Action

(None).

## Syslog Message

Aggressive timeout flush ended on disk %s S/N %s.

## Parameters

**disk\_name** (STRING): Name of the disk.

**serial\_num** (STRING): Serial number of the disk.

## disk.timeout.flush.start

### Severity

NOTICE

### Description

This message occurs when an I/O operation is taking longer than expected and aggressive timeout I/Os are flushed from the disk driver.

### Corrective Action

(None).

## Syslog Message

Aggressive timeout flush started on disk %s S/N %s. Details: %s.

## Parameters

**disk\_name** (STRING): Name of the disk.

**serial\_num** (STRING): Serial number of the disk.

**details** (STRING): Additional information.

# disk.toomanynvmespindles events

## disk.tooManyNvmeSpindles

### Severity

ALERT

### Description

This message occurs when Data ONTAP® finds more NVMe disks present on a platform than the platform actually supports. Additional disks above the support number are ignored. Correct this issue before rebooting the system.

### Corrective Action

Too many NVMe disks have been installed. Reconfigure the system to the maximum number of NVMe disks supported.

### Syslog Message

Reached the maximum number of NVMe disks that are supported on this system of %d NVMe disks. All additional NVMe disks will be ignored. This problem must be resolved before the system is rebooted.

### Parameters

**maxSpindle** (INT): Maximum number of NVMe disks supported on this platform.

# disk.toomanypartitioneddisk events

## disk.tooManyPartitionedDisks

### Severity

ALERT

### Description

This message occurs when Data ONTAP® finds more partitioned disks present on a platform than the platform supports. The system ignores any partitioned disks above the limit, and they will be unusable.

### Corrective Action

Too many partitioned disks have been installed. Reconfigure the system to the maximum number of partitioned disks supported.

### Syslog Message

Ignoring partitioned disk %s because the maximum limit of %d partitioned disks supported on this system has been reached.

### Parameters

**diskName** (STRING): Name of the disk.

**maxPartitionedDisks** (INT): Maximum number of partitioned disks supported on this platform.

# disk.toomanyreassignments events

## disk.tooManyReassignments

### Severity

NOTICE

### Description

This message occurs when a reassignment fails due to table overflow.

### Corrective Action

(None).

### Syslog Message

Disk %s: sector %llu was NOT reassigned, too many reassignments.

### Parameters

**diskName** (STRING): Name of the device.

**error** (LONGINT): Bad sector.

## disk.toomanyspindles events

### disk.tooManySpindles

### Severity

ALERT

### Description

This message occurs when Data ONTAP® finds more disks present on a platform than the platform actually supports. Additional disks above the support number are ignored. Correct this issue before rebooting the system.

### Corrective Action

Too many disks have been installed. Reconfigure the system to the maximum number of disks supported.

### Syslog Message

Found %d disks, but maximum of %d disks are supported on this system. All additional disks will be ignored. This problem must be resolved before the system is rebooted.

### Parameters

**actualSpindle** (INT): Actual number of disks configured.

**maxSpindle** (INT): Maximum number of disks supported on this platform.

## disk.uid events

### disk.uid.changed

### Severity

NOTICE

## Description

This message occurs when a SCSI (Small Computer System Interface) INQUIRY command done during a disk drive initialization sequence finds different data in the INQUIRY response for page 0x83. The unique id (UID) derived from this information has changed unexpectedly.

## Corrective Action

Locate the failed disk drive by looking for a disk drive whose amber LED fault light is lit. Replace the disk drive.

## Syslog Message

Disk %s has unexpectedly changed its unique ID. vendor:%s product:%s serial\_number:%s

## Parameters

**diskName** (STRING): Name of the disk that had the UID change.  
**disk\_vendor\_id** (STRING): Vendor name of the disk that had the UID change.  
**disk\_product\_id** (STRING): Product name of the disk that had the UID change.  
**disk\_serialno** (STRING): Serial number of the disk that had the UID change.

# disk.unabletorewritedata events

## disk.unableToRewriteData

### Severity

INFORMATIONAL

### Description

This message occurs when Data ONTAP® receives too many errors reading from a particular location on an ATA disk drive, but its attempt to rewrite the data to the disk drive fails.

### Corrective Action

(None).

### Syslog Message

Difficulty reading disk drive %s with the command %s; could not rewrite the data.

### Parameters

**diskName** (STRING): Name of the disk drive with the problem.  
**command** (STRING): SCSI command that reported a problem.

# disk.vendorspecificcode events

## disk.vendorSpecificCode

### Severity

ERROR

### Description

This message occurs when a disk returns a vendor-specific error.

## Corrective Action

Replace the disk.

## Syslog Message

Drive %s has failed with a vendor specific error: class 0x%x segment 0x%x sense\_key 0x%x info 0x%x 0x%x 0x%x 0x%x length 0x%x resv\_1 0x%x sense\_code %d sense\_code\_qualifier %d fru\_failed 0x%x flags 0x%x field\_pointer 0x%x.

## Parameters

**diskName** (STRING): Name of the disk.  
**class** (INHEX): Sense class.  
**segment** (INHEX): Segment.  
**sense\_key** (INHEX): Actual sense key.  
**info0** (INHEX): Vendor-specific information[0].  
**info1** (INHEX): Vendor-specific information[1].  
**info2** (INHEX): Vendor-specific information[2].  
**info3** (INHEX): Vendor-specific information[3].  
**len** (INHEX): Length.  
**resv** (INHEX): Resv field.  
**sense\_code** (INT): Actual sense code.  
**qualifier** (INT): Actual sense code qualifier.  
**fru\_failed** (INHEX): Actual FRU code.  
**flags** (INHEX): Flags of the error.  
**field\_ptr** (INHEX): Sense field pointer.

# disk.writeverifyerror events

## disk.writeVerifyError

### Severity

ERROR

### Description

This message occurs when a write verify error occurs.

### Corrective Action

(None).

### Syslog Message

WRITE VERIFY on disk %s received an unexpected status %d.

### Parameters

**diskName** (STRING): Name of the device.  
**error** (INT): Returned error code



# diskown events

## diskown.aashelf events

### diskown.AAShelf.InterPools

#### Severity

ERROR

#### Description

This message occurs when the controller cannot automatically assign a disk because disks on the shelf and disks on the loop have the same owner but belong to different pools.

#### Corrective Action

1) Find all of the disks belonging to the same shelf, loop and owners as disk %s (S/N %s) by using the "storage disk show" command with the "-fields shelf, owner, dr-home, pool" option. 2) Assign all of the disks with the same owners on this shelf and other shelves to the same pool. Alternatively, manually assign all unowned disks belonging to this shelf.

#### Syslog Message

The controller could not automatically assign disk %s (S/N %s) and any other unassigned disks in the same shelf because the disks are shared across multiple pools within the same loop.

#### Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

### diskown.AAShelf.MoreShelf

#### Severity

ERROR

#### Description

This message occurs when more than the supported number of shelves are connected to a storage stack.

#### Corrective Action

1) Check which adapter has more than the supported number of shelves attached by using the "run local storage show shelf" command. 2) Reconfigure the shelves so that the adapter is not connected to more than the supported number of shelves.

#### Syslog Message

(None).

#### Parameters

(None).

### diskown.AAShelf.MultiOwners

**Severity**

ERROR

**Description**

This message occurs when the controller cannot automatically assign disks because the disks on the shelf are owned by multiple systems.

**Corrective Action**

1) Find all of the disks belonging to the same shelf as disk %s (S/N %s) by using the "storage disk show" command. 2) View the list of systems by using the "run local disk show -v" command. 3) Either assign all of the disks on this shelf to a single system, or assign all unowned disks belonging to this shelf.

**Syslog Message**

The controller could not automatically assign disk %s (S/N %s) because disks on the shelf are owned by multiple systems. Because of this problem, the controller also could not assign any unowned disks on this shelf.

**Parameters**

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

**diskown.AAShelf.MultiPools****Severity**

ERROR

**Description**

This message occurs when the controller cannot automatically assign a disk because the disks on the shelf are shared across multiple pools.

**Corrective Action**

1) Find all of the disks belonging to the same shelf as disk %s (S/N %s) by using the "storage disk show" command. 2) View the list of pools by using the "run local disk show -v" command. 3) Either assign all of the disks on this shelf to a single pool, or assign all unowned disks belonging to this shelf.

**Syslog Message**

The controller could not automatically assign disk %s (S/N %s) and any other unassigned disks in the same shelf because the disks are shared across multiple pools.

**Parameters**

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

**diskown.AAShelf.NoMemory****Severity**

ALERT

**Description**

This message occurs when the node cannot allocate memory for shelf-level automatic assignment.

### Corrective Action

Initiate the automatic disk assign by running the "storage disk assign -auto true -node node-name" CLI command. If the out-of-memory condition persists, perform a manual disk assign.

### Syslog Message

Memory not allocated for shelf-level automatic assignment.

### Parameters

(None).

## diskown.AAShelf.NoOwner

### Severity

ERROR

### Description

This message occurs when the controller cannot automatically assign a disk because none of the disks on the shelf are owned by any system.

### Corrective Action

1) Find all of the disks belonging to the same shelf as disk %s (S/N %s) by using the "storage disk show" command. 2) View the list of systems by using the "run local disk show -v" command. 3) Either assign all of unowned disks in this shelf or assign one of the disks on this shelf by using the "storage disk assign" command.

### Syslog Message

The controller could not automatically assign disk %s (S/N %s) (and other unassigned disks in the same shelf, if any) because no disks on the shelf are owned by any system.

### Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

## diskown.assign events

### diskown.assign.splitOwner

### Severity

INFORMATIONAL

### Description

This message occurs when the "storage disk assign" command fail to assign a disk to a different owner from the other disks within the carrier. This is because the "storage disk assign" command with the -f option must be used in such cases.

### Corrective Action

Use the "storage disk assign" command with the -f option to assign the disk to a different owner from the other disks within the carrier.

## Syslog Message

Assigning ownership failed for disk %s. Use the "storage disk assign" command with the -f option to assign the disk to a different owner from the other disks within the carrier.

## Parameters

**diskname** (STRING): Name of the disk.

# diskown.assignexceeded events

## diskown.AssignExceeded

### Severity

ERROR

### Description

This message occurs when an attempt is made to assign more disks than are supported by this node.

### Corrective Action

Review the system capacity limits and the assigned disks on this system.

## Syslog Message

The maximum number of disks that can be assigned to this node has been exceeded. Attempted to assign %d, maximum supported: %d.

## Parameters

**attempted** (INT): Number of disks attempted to assign.

**supported** (INT): Number of disks supported.

# diskown.autoassign events

## diskown.AutoAssign.MultipleOwners

### Severity

ERROR

### Description

This message occurs when automatic assigning of a disk fails because the disks on the loop are owned by multiple systems.

### Corrective Action

Use the "storage disk show" command to view the list of systems, and then assign all of the disks on this loop to a single system. Or use the "storage disk assign" command to manually assign this disk.

## Syslog Message

Automatic assigning failed for disk %s (S/N %s) because the disks on the loop are owned by multiple systems. Automatic assigning failed for all unowned disks on this loop.

## Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

## diskown.AutoAssign.MultiplePools

### Severity

ERROR

### Description

This message occurs when automatic assigning of a disk fails because the disks on the loop are shared across multiple pools.

### Corrective Action

Use the "storage disk assign -pool" command to assign all of the disks on this loop to a single pool. Or use the "storage disk assign" command to manually assign this disk.

### Syslog Message

Automatic assigning failed for disk %s (S/N %s) because the disks on the loop are shared across multiple pools. Automatic assigning failed for all unowned disks on this loop.

### Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

## diskown.AutoAssign.NoOwner

### Severity

ERROR

### Description

This message occurs when automatic assigning of a disk fails because none of the disks on the loop are owned by any system.

### Corrective Action

Use the "run local disk show -v" command to get the list of systems and unowned disks. Manually assign one of the disks on this loop by using the "storage disk assign" command.

### Syslog Message

Automatic assigning failed for disk %s (S/N %s) because none of the disks on the loop are owned by any system. Automatic assigning failed for all unowned disks on this loop.

### Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

## diskown.autoassignproblem events

### diskown.AutoAssignProblem

### Severity

ERROR

## Description

This message occurs when Data ONTAP® attempts to assign a disk automatically, but the request fails. The failure could be due to a disk error or a user simultaneously assigning the disk to a storage system.

## Corrective Action

Use the "disk show -v" nodeshell command to determine the status of the disk. Use the "storage disk assign" command to assign the disk.

## Syslog Message

Automatic assigning failed for disk %s (S/N %s). Reason: %s

## Parameters

**diskname** (STRING): Name of the disk.  
**serialno** (STRING): Serial number of the disk.  
**reason** (STRING): Reason for the failure.

# diskown.bayctrlr events

## diskown.bayctrlr

### Severity

NOTICE

### Description

This message occurs if the controller address is determined when the bay policy override option has been set. The controller address "A/B" is used to automatically assign the disks in odd/even bays.

### Corrective Action

(None).

### Syslog Message

Determined controller address %s based on shelf ID %d and slot number %d.

### Parameters

**ctrlr** (STRING): Controller address.  
**shelfid** (INT): ID of the shelf chosen to determine the controller address.  
**slot** (INT): Slot number of the shelf chosen to determine the controller address.

## diskown.bayctrlr.unknown

### Severity

NOTICE

### Description

This message occurs when disks in a shelf are not assigned to the current node.

### Corrective Action

(None).

## Syslog Message

Cannot make determination of controller address.

## Parameters

(None).

# diskown.changingowner events

## diskown.changingOwner

### Severity

NOTICE

### Description

This message occurs when the ownership of a disk is being changed.

### Corrective Action

(None).

## Syslog Message

The ownership of disk %s (S/N %s) is being changed from node "%s" (ID: %u, home ID: %u, DR home ID: %u) to node "%s" (ID: %u, home ID: %u, DR home ID: %u).

## Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

**oldownername** (STRING): Name of the former owner node.

**oldownerid** (LONGINT): System ID of the former owner node.

**oldhomeownerid** (LONGINT): System ID of the former home node.

**olddrhomeownerid** (LONGINT): System ID of the former disaster recovery home node.

**newownername** (STRING): Name of the new owner node.

**newownerid** (LONGINT): System ID of the new owner node.

**newhomeownerid** (LONGINT): System ID of the new home node.

**newdrhomeownerid** (LONGINT): System ID of the new disaster recovery home node.

**thread** (STRING): Name of the thread that issued the ownership change request.

**APIname** (STRING): Name of the API that issued the ownership change request.

# diskown.cleaningownership events

## diskown.cleaningOwnership

### Severity

INFORMATIONAL

### Description

This message occurs when a "storage disk assign -f" command is run for an SES drive and the system discovers that the newly assigned owner already owns an SES drive on the same loop and shelf as the supplied drive. The assign is completed because it is a forced assign, but the error message is generated as an alert.

### Corrective Action

(None).

### Syslog Message

Disk %s (S/N %s) contains software-based disk ownership information that indicates it was previously in use by system ID %u. The ownership information is being cleared.

### Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

**ownerid** (INT): System ID of the system that owned the disk.

## diskown.completion events

### diskown.completion.timer.reEnabled

#### Severity

INFORMATIONAL

#### Description

This message occurs when the COMPLETION events from sanown layer remain after an aggregate is migrated from a source to a destination. Typically the events are enabled automatically, and if that doesn't happen within a certain amount of time, the system reenables the COMPLETION events and logs this message.

#### Corrective Action

(None).

#### Syslog Message

Sanown completion event timer reenabled after a giveback.

#### Parameters

(None).

## diskown.delayed events

### diskown.delayed.completionEvent

#### Severity

INFORMATIONAL

#### Description

This message occurs when the software disk ownership module completes the process of updating its view of the disks attached to the system but delays notifying other layers of the completion of this process. Typically, this happens when an SFO aggregate giveback is in progress on the HA partner.

#### Corrective Action

(None).



### **Syslog Message**

Software disk ownership module temporarily delayed notification that it has completed updating its view of the disks.

### **Parameters**

(None).

## **diskown.disk events**

### **diskown.disk.already.deleted**

#### **Severity**

ERROR

#### **Description**

This message occurs during an attempt to assign a disk that is being deleted. This can happen if the disk is deleted at the storage layer while disk assignment is in progress.

#### **Corrective Action**

Run the "run local disk show -v" command to see the list of available disks, and then retry the operation using one of the listed disks.

### **Syslog Message**

Assign request failed for disk %s because it is being deleted.

### **Parameters**

**disk\_name** (STRING): Name of the disk that is being deleted.

## **diskown.diskhaswronginfo events**

### **diskown.diskHasWrongInfo**

#### **Severity**

ERROR

#### **Description**

This message occurs when the disk ownership information on a disk does not belong to the disk that it was read from. This could be a result of a read going to the wrong disk or a prior write intended for some other disk getting written to the disk.

#### **Corrective Action**

On the controller that owns the disk, use the "storage disk assign" command to rewrite the ownership information to the disk.

### **Syslog Message**

Disk %s (S/N %s) has ownership information for disk %s (S/N %s) on it.

### **Parameters**

**diskname** (STRING): Name of the disk from which the ownership information was read.

**serialno** (STRING): Serial number of the disk from which the ownership information was read.

**infoForDiskname** (STRING): Name of the disk that the ownership information is for.

**infoForSerialNo** (STRING): Serial number of the disk that the ownership information is for.

## diskown.disknotsup events

### diskown.diskNotSup

#### Severity

INFORMATIONAL

#### Description

This message occurs when newly inserted disks are only supported with software disk ownership enabled, but it is not enabled.

#### Corrective Action

Halt the system by using the "halt" command, and then turn on software disk ownership by using the "storage disk assign" command. After you assign the disks, the system can recognize them.

#### Syslog Message

Newly added disk %s (and others) is only supported with software disk ownership enabled.

#### Parameters

**disk** (STRING): Name of the disk.

## diskown.diskwasstolen events

### diskown.diskWasStolen

#### Severity

NOTICE

#### Description

This event is issued when the ownership of a disk owned by the system is changed such that the disk is no longer owned by the system.

#### Corrective Action

(None).

#### Syslog Message

Disk %s (S/N %s) has had its ownership changed so that it is no longer owned by this system. This can lead to a system panic if the disk is a filesystem disk.

#### Parameters

**diskname** (STRING): The name of the disk

**serialno** (STRING): The serial number of the disk

## diskown.drawer events

## **diskown.drawer.assignStatus**

### **Severity**

NOTICE

### **Description**

This message occurs when drawer-based drive assignment for shelves with drawers is enabled or disabled.

### **Corrective Action**

(None).

### **Syslog Message**

Drawer-based automatic drive assignment for shelf "%d" is "%s".

### **Parameters**

**shelf\_id** (INT): Shelf identifier.

**state** (STRING): Drawer-based automatic drive assignment state..

## **diskown.drawer.invalidShlf**

### **Severity**

ERROR

### **Description**

This message occurs when the controller cannot automatically assign disks because the shelf is either offline, or temporarily unavailable, and therefore cannot be queried for the information that is required to properly assign the disks to the controller. This can also occur when multiple shelves in a cluster have the same shelf ID. Automatic disk assignment fails for the disks in this shelf until the shelf is back online or any shelf ID conflicts within a cluster are resolved.

### **Corrective Action**

Use the "storage shelf show" command to verify that the specified shelf is online and to verify connectivity from the controller to all of the paths to the shelf. If multiple shelves in a cluster have the same shelf ID, resolve that and ensure that shelves in that cluster have different shelf IDs.

### **Syslog Message**

The system cannot automatically assign disks for shelf "%d" because the shelf is in an invalid state.

### **Parameters**

**shelf\_id** (INT): Shelf identifier.

## **diskown.drawer.MultiOwners**

### **Severity**

ERROR

### **Description**

This message occurs when the system cannot automatically assign disks because some of the disks are assigned to multiple systems in a conflicting manner.

### Corrective Action

1) Find all of the disks belonging to the specified shelf by using the "storage disk show" command. 2) View the list of systems by using the "disk show -v" command. 3) Ensure that disks assigned on lower half bays of each drawer are owned by a single owner, and ensure that disks assigned on higher half bays of each drawer are owned by a different single owner. Disks assigned on lower half bays of each drawer must have a different owner than disks assigned on higher half bays.

### Syslog Message

The system cannot automatically assign disks because the disks in shelf "%d" are assigned to multiple systems in a conflicting manner.

### Parameters

**shelf\_id** (INT): Shelf identifier.

## diskown.duplicatessoftentries events

### diskown.duplicateSoftEntries

#### Severity

NOTICE

#### Description

This event is issued when we find duplicate entries in the sanown filer table.

#### Corrective Action

(None).

#### Syslog Message

%s

#### Parameters

**ems\_txt** (STRING): The warning message

## diskown.errorduringio events

### diskown.errorDuringIO

#### Severity

ERROR

#### Description

This event is issued when an I/O error occurs during a sanown disk operation.

#### Corrective Action

(None).

#### Syslog Message

error %d (%s) on disk %s (S/N %s) while %s

## Parameters

**errno** (INT): The error number that occurred  
**errortext** (STRING): The error description as text  
**diskname** (STRING): The name of the disk  
**serialno** (STRING): The serial number of the disk  
**details** (STRING): The operation that was going on

# diskown.errorreadingownership events

## diskown.errorReadingOwnership

### Severity

NOTICE

### Description

This event is issued when we get an error while reading all the DOWNS at once. The recovery is to read them individually.

### Corrective Action

(None).

### Syslog Message

error %d (%s) while reading ownership on disk %s (S/N %s)

## Parameters

**errno** (INT): The error that occurred  
**errortext** (STRING): The error description as text  
**diskname** (STRING): The name of the disk  
**serialno** (STRING): The serial number of the disk

# diskown.generated events

## diskown.generated.completionEvent

### Severity

INFORMATIONAL

### Description

This message occurs when the software disk ownership module notifies other layers that it has completed the process of updating its view of the disks attached to the system. These notifications were previously disabled for a short time.

### Corrective Action

(None).

### Syslog Message

Software disk ownership module notified other layers that it has completed updating its view of the disks.

## Parameters

(None).

# diskown.hlfshlf events

## diskown.hlfShlf.assignStatus

### Severity

NOTICE

### Description

This message occurs when half-shelf based drive assignment is enabled or disabled.

### Corrective Action

(None).

### Syslog Message

Half-shelf based automatic drive assignment for shelf "%d" is "%s".

### Parameters

**shelf\_id** (INT): Shelf identifier.

**state** (STRING): Half-shelf based automatic drive assignment state.

## diskown.hlfShlf.invalidShlf

### Severity

NOTICE

### Description

This message occurs when the controller cannot automatically assign disks to a shelf using half shelf-based disk assignment, either because the shelf is offline, or is temporarily unavailable and therefore cannot be queried for the information required to properly divide the disks for assignment to the controller. This can also occur when multiple shelves in a cluster have the same shelf ID. Automatic half shelf-based disk assignment is bypassed for disks in this shelf until the shelf is back online or any shelf ID conflicts within a cluster are resolved.

### Corrective Action

Use the "storage shelf show" command to ensure that the specified shelf is online and to ensure connectivity from the controller to all paths to the shelf. If there are multiple shelves in a cluster with the same same shelf ID, resolve that and ensure that shelves in that cluster have different shelf IDs.

### Syslog Message

Half shelf based disk assignment will be skipped for shelf "%d" because the shelf is in invalid state.

### Parameters

**shelf\_id** (INT): Shelf identifier.

## diskown.hlfShlf.MultiOwners

## Severity

NOTICE

## Description

This message occurs when the controller cannot automatically assign disks using half shelf based disk assignment because disks in a shelf are assigned to multiple systems in a conflicting manner.

## Corrective Action

1) Find all of the disks belonging to the shelf specified by using the "storage disk show" command. 2) View the list of systems by using the "run local disk show -v" command. 3) Ensure that disks assigned on lower half bays of the shelf are owned by a single owner. Similarly, ensure that disks assigned on higher half bays of that shelf are owned by a single owner. Ensure that disks assigned on lower half bays have a different owner than disks assigned on higher half bays.

## Syslog Message

The controller cannot automatically assign disks using half shelf based disk assignment because disks in shelf "%d" are assigned to multiple systems in a conflicting manner.

## Parameters

**shelf\_id** (INT): Shelf identifier.

# diskown.ilv events

## diskown.ilv.hShlf.invShlf

## Severity

ERROR

## Description

This message occurs when the controller cannot automatically assign drives to a shelf using interleaved-half-shelf-based drive assignment, either because the shelf is offline, or is temporarily unavailable and therefore cannot be queried for the information required to properly divide the drives for assignment to the controller. This can also occur when multiple shelves in a cluster have the same shelf ID. Automatic interleaved-half-shelf-based drive assignment is bypassed for drives in this shelf until the shelf is back online or any shelf ID conflicts within a cluster are resolved.

## Corrective Action

Use the "storage shelf show" command to ensure that the specified shelf is online and to ensure connectivity from the controller to all paths to the shelf. Ensure that each shelf within the cluster has a unique shelfID.

## Syslog Message

Interleaved-half-shelf-based drive assignment will be skipped for shelf "%d" because the shelf is in invalid state.

## Parameters

**shelf\_id** (INT): Shelf identifier.

## diskown.ilv.hShlf.mOwners

**Severity**

ERROR

**Description**

This message occurs when the controller cannot automatically assign drives using interleaved-half-shelf based drive assignment because drives in a shelf are assigned to multiple systems in a conflicting manner.

**Corrective Action**

1) Find all of the drives belonging to the shelf specified by using the "storage disk show" command. 2) View the list of systems by using the "run local disk show -v" command. 3) Ensure that drives assigned on first and third quarter bays of the shelf are owned by a single owner. Similarly, ensure that drives assigned on second and fourth quarter bays of that shelf are owned by a single owner. Ensure that drives assigned on first and third quarter bays have a different owner than drives assigned on second and fourth quarter bays.

**Syslog Message**

The controller cannot automatically assign drives using interleaved-half-shelf based drive assignment because drives in shelf "%d" are assigned to multiple systems in a conflicting manner.

**Parameters**

**shelf\_id** (INT): Shelf identifier.

**diskown.ilv.hShlf.status****Severity**

NOTICE

**Description**

This message occurs when interleaved-half-shelf based drive assignment is either enabled or disabled.

**Corrective Action**

(None).

**Syslog Message**

Interleaved-half-shelf based automatic drive assignment for shelf %d is "%s".

**Parameters**

**shelf\_id** (INT): Shelf identifier.

**state** (STRING): Interleaved-half-shelf based automatic drive assignment state.

**diskown.incorrectownername events****diskown.IncorrectOwnerName****Severity**

INFORMATIONAL

**Description**

This event is issued when we find a system finds a disk that it owns, but which has a different owner name than the system name.



### Corrective Action

(None).

### Syslog Message

Disk %s has incorrect owner name information (%s). It will be changed to match the system name (%s).

### Parameters

**diskname** (STRING): The name of the disk for which this was observed

**wrong\_ownership** (STRING): The owner name, which is incorrect, that the disk currently has

**system\_name** (STRING): The system name that should be the owner name for the disk

## diskown.isbeingenabled events

### diskown.isBeingEnabled

#### Severity

INFORMATIONAL

#### Description

This event is issued when the system is being converted from non-software based disk ownership to software based disk ownership

#### Corrective Action

(None).

#### Syslog Message

System is being changed to use software based disk ownership

#### Parameters

(None).

## diskown.isenabled events

### diskown.isEnabled

#### Severity

INFORMATIONAL

#### Description

This event is issued when enable sanown.

#### Corrective Action

(None).

#### Syslog Message

software ownership has been enabled for this system

#### Parameters

(None).

# diskown.mcc events

## diskown.mcc.AssignExceeded

### Severity

ERROR

### Description

This message occurs when the number of drives assigned in the MetroCluster(tm) configuration exceeds the supported limit.

### Corrective Action

Review the system capacity limits and the assigned drives on this system. If possible, unassign the unused array LUNs. If you need further assistance, contact NetApp technical support.

### Syslog Message

The maximum number of drives that can be assigned in the configuration has exceeded. Total physical drive count: %d, maximum supported: %d.

### Parameters

**total\_disk\_count** (INT): Total number of drives assigned in the configuration.

**supported** (INT): Number of drives supported.

# diskown.mccip events

## diskown.mccip.assignStatus

### Severity

NOTICE

### Description

This message occurs when MetroCluster(tm) quarter-shelf based drive assignment is enabled or disabled. It is enabled when the system is a member of an AFF HA pair in a 4-node MetroCluster IP configuration.

### Corrective Action

(None).

### Syslog Message

MetroCluster quarter-shelf based automatic drive assignment for shelf "%d" is "%s".

### Parameters

**shelf\_id** (INT): Shelf identifier.

**state** (STRING): MetroCluster quarter-shelf based automatic drive assignment state.

## diskown.mccip.invalidShlf

### Severity

ERROR

## Description

This message occurs when the controller cannot automatically assign disks because the shelf is either offline, or temporarily unavailable and therefore cannot be queried for the information that is required to properly assign the disks to the controller. This can also occur when multiple shelves in a cluster have the same shelf ID. Automatic disk assignment fails for the disks in this shelf until the shelf is back online or any shelf ID conflicts within a cluster are resolved.

## Corrective Action

Use the "storage shelf show" command to verify that the specified shelf is online and to verify connectivity from the controller to all of the paths to the shelf. If multiple shelves in a cluster have the same shelf ID, resolve that and ensure that shelves in that cluster have different shelf IDs.

## Syslog Message

The system cannot automatically assign disks for shelf "%d" because the shelf is in an invalid state.

## Parameters

**shelf\_id** (INT): Shelf identifier.

## diskown.mccip.MultiOwners

### Severity

ERROR

## Description

This message occurs when the system cannot automatically assign disks because some of the disks are assigned to multiple systems in a conflicting manner.

## Corrective Action

1) Find all of the disks belonging to the specified shelf by using the "storage disk show" command. 2) View the list of systems by using the "disk show -v" command. 3) Ensure that the disks on each shelf are assigned such that lower half bays are assigned to the local high-availability (HA) pair and the higher half bays are assigned to the remote HA pair. Additionally within each half shelf, ensure that the disks are assigned equally to the HA pair such that the lower quarter bays are assigned to one node and the higher quarter bays are assigned to the node's HA partner.

## Syslog Message

The system cannot automatically assign disks because the disks in shelf "%d" are assigned to multiple systems in a conflicting manner.

## Parameters

**shelf\_id** (INT): Shelf identifier.

## diskown.ownerchangefailed events

### diskown.ownerChangeFailed

### Severity

ERROR

## Description

This message occurs when a disk ownership change request fails.

## Corrective Action

Verify that there are no disk connectivity issues, that disks are not in a failed state, and that the procedures for executing the "storage disk assign" command are being followed.

## Syslog Message

Disk ownership change request failed on disk %s (S/N %s). Reason: %s.

## Parameters

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

**status** (STRING): Description of the error in the ownership change request.

**thread** (STRING): Name of the thread that issued the ownership change request.

**APIname** (STRING): Name of the API that issued the ownership change request.

# diskown.ownernameupdatefailed events

## diskown.OwnerNameUpdateFailed

### Severity

NOTICE

### Description

This event is issued when we have to rewrite a DOWN because it is out of date wrt the other DOWNS.

### Corrective Action

(None).

### Syslog Message

owner name update failed for disk %s (S/N %s)

### Parameters

**diskname** (STRING): The name of the disk

**serialno** (STRING): The serial number of the disk

# diskown.ownerreservationmismatch events

## diskown.ownerReservationMismatch

### Severity

ERROR

### Description

This message occurs when the system discovers a disk that is supposed to be owned by one node, but the persistent reservation is owned by another node.

## Corrective Action

1) Find the disk causing the conflict from EMS messages. 2) If the disk is in the spare pool, go to step 5. 3) Run the "storage aggr status -r" command to find the aggregate containing the disk. 4) If the aggregate is online, bring it offline by using the "storage aggr modify -aggr aggr\_name -state offline" command. 5) Run the "storage disk show -fields owner-id,reserver-id" command to identify the reservation owner of the disk. 6) Run the "storage disk assign disk\_name -s partner\_sys\_id -f" command to perform the disk assignment change from the node that owns the reservation for the disk.

## Syslog Message

disk %s (S/N %s) is supposed to be owned by this node but has a persistent reservation placed by node %s (ID %u).

## Parameters

**diskname** (STRING): Name of the disk.  
**serialno** (STRING): Serial number of the disk.  
**resvownername** (STRING): Name of the system owning the reservation.  
**resvowner** (INT): System ID of the node owning the reservation.

# diskown.raid events

## diskown.RAID.Ownershipchange

### Severity

ERROR

### Description

This message occurs when a disk assign or disk remove ownership request is rejected at the RAID layer.

### Corrective Action

Review the error message and verify that the request was not for a file system disk.

### Syslog Message

%s

### Parameters

**err\_msg** (STRING): Reason the request failed.

# diskown.reassignfailed events

## diskown.ReassignFailed

### Severity

ERROR

### Description

This message occurs when a disk reassign request fails.

### Corrective Action

Verify that there are no disk connectivity issues and that the procedures for executing the 'storage disk reassign' command are being followed.

### Syslog Message

Disk reassign request failed. Reason: %s.

### Parameters

**reason** (STRING): Reason for failure.

## diskown.releasingmismatchedreservation events

### diskown.releasingMismatchedReservation

#### Severity

NOTICE

#### Description

This event is issued when we release the reservation because the DOWN says the disk is owned by someone else.

#### Corrective Action

(None).

#### Syslog Message

%s is releasing reservation on disk %s (S/N %s) as it is owned by %s (ID %u)

#### Parameters

**originator** (STRING): The internal component of diskown issuing the event

**diskname** (STRING): The name of the disk

**serialno** (STRING): The serial number of the disk

**filername** (STRING): The name of the filer that owns the disk

**systemid** (INT): The system id of the filer that owns the disk

## diskown.replayingownership events

### diskown.replayingOwnership

#### Severity

INFORMATIONAL

#### Description

This event occurs during a boot or takeover when a disk ownership request is replayed from NVRAM.

#### Corrective Action

(None).

#### Syslog Message

Replaying ownership request for %d disks

#### Parameters

**diskcount** (INT): The number of disks in the ownership request

## diskown.replayingOwnership.Nv

### Severity

ERROR

### Description

This message occurs during takeover, when a disk ownership request is replayed from NVRAM but NVRAM corruption is detected.

### Corrective Action

1) If you changed the ownership of disks on the partner in the last five minutes prior to takeover, verify their ownerships were stored by using the "storage disk show" command. 2) If you see this message often or in other modules, the partner's NVRAM card might be faulty and might need to be replaced.

### Syslog Message

Failed to replay disk ownership because NVRAM corruption was detected.

### Parameters

(None).

## diskown.rescan events

### diskown.rescan.notConnected

### Severity

ERROR

### Description

This message occurs when the high-availability (HA) messaging layer fails to connect to the specified node for disk ownership updates.

### Corrective Action

Run the command "run local disk show -v" from the node specified in the node parameter to correct disk ownership.

### Syslog Message

HA messaging layer on the local node failed to connect to node %s for disk ownership updates.

### Parameters

**node** (STRING): Name of the node where disk ownership might be outdated.

**status** (INT): Error number returned by the HA messaging layer from the HA message error space.

**scope** (INT): Internal scope identifier for the module that returned the error.

**scope\_err** (INT): Internal error number returned by the module.

### diskown.rescan.notSent

### Severity

ERROR

## Description

This message occurs when the high-availability (HA) messaging layer fails to send a disk ownership update message.

## Corrective Action

Run the command "run local disk show -v" from the node specified in the node parameter to correct disk ownership.

## Syslog Message

HA messaging layer on the local node failed to send a disk ownership update message to node %s.

## Parameters

**node** (STRING): Name of the node where disk ownership might be outdated.

**status** (INT): Error number returned by the HA messaging layer from the HA message error space.

**scope** (INT): Internal scope identifier for the module that returned the error.

**scope\_err** (INT): Internal error number returned by the module.

# diskown.rescanmessagefailed events

## diskown.RescanMessageFailed

### Severity

ERROR

### Description

This message occurs while transferring the ownership of a disk to another node, when a rescan message is sent but does not go through.

### Corrective Action

Use the "disk show" command on the nodeshell of %s for the node to rescan the disks which changed ownership.

### Syslog Message

Could not send rescan message to %s.

### Parameters

**hostname** (STRING): Name of the host that the message was sent to.

# diskown.rescanwait events

## diskown.rescanWait

### Severity

INFORMATIONAL

### Description

This message occurs when the disk ownership rescan code waits.



**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**status** (STRING): Description on the type of wait.

## diskown.rewriteowner events

**diskown.rewriteOwner****Severity**

NOTICE

**Description**

This event is issued when we have to rewrite a DOWN because it is out of date wrt the other DOWNS.

**Corrective Action**

(None).

**Syslog Message**

rewriting ownership information for disk %s (S/N %s)

**Parameters**

**diskname** (STRING): The name of the disk

**serialno** (STRING): The serial number of the disk

## diskown.shareddisk events

**diskown.sharedDsk.nonHACfg****Severity**

ERROR

**Description**

This message occurs when a disk is shared by multiple nodes in a non-HA or stand-alone configuration. If the procedure to convert an HA configuration to a non-HA configuration is not followed properly, the system may end up either having disks whose partitions are owned by another node, or owning partitions on non-local disks.

**Corrective Action**

Refer to the High-Availability Configuration Guide.

**Syslog Message**

Either some local disks have partitions that are not owned by the local node, or the local node owns partitions on non-local disks.

## Parameters

(None).

# diskown.shlf events

## diskown.Shlf.assignStatus

### Severity

NOTICE

### Description

This message occurs when shelf granular-based drive assignment is enabled or disabled. This is currently applicable for an MetroCluster(tm) IP configurations.

### Corrective Action

(None).

### Syslog Message

Shelf-based automatic drive assignment for shelf "%d" is "%s".

### Parameters

**shelf\_id** (INT): Shelf identifier.

**state** (STRING): Shelf-based automatic drive assignment state.

## diskown.Shlf.MultiOwners

### Severity

ERROR

### Description

This message occurs when the system cannot automatically assign disks because some of the disks are assigned to multiple nodes in a conflicting manner. Until the disk assignments are corrected, the specified shelf will be excluded during auto assignment, and unowned disks on the shelf will be skipped.

### Corrective Action

1) Find all of the disks belonging to the specified shelf by using the "storage disk show -fields owner,shelf" command. 2) Ensure that the disks on the shelf are assigned to a single system by manually assigning disks in the shelf to a single owner.

### Syslog Message

The system cannot automatically assign disks because the disks in shelf %d are assigned to multiple nodes in a conflicting manner.

### Parameters

**shelf\_id** (INT): Shelf identifier.

# diskown.split events

## diskown.split.shelf.assignStatus

### Severity

NOTICE

### Description

This message occurs when any type of split-shelf automatic drive assignment is disabled.

### Corrective Action

(None).

### Syslog Message

Split-shelf based automatic drive assignment is "%s".

### Parameters

**state** (STRING): All split-shelf based drive assignment states.

## diskown.statusreadownership events

### diskown.statusReadOwnership

### Severity

INFORMATIONAL

### Description

This message occurs when an event is logged to track that an error was encountered while reading all data ownership blocks at once. Data ONTAP® will attempt to recover by reading each disk ownership block individually.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**errorno** (INT): Error that occurred.

**errortext** (STRING): Error description as text.

**diskname** (STRING): Name of the disk.

**serialno** (STRING): Serial number of the disk.

## diskown.stealingdisk events

### diskown.StealingDisk

### Severity

NOTICE

### Description

This event is issued when a user reassigns a disk that is owned and has a reservation from another system.

## Corrective Action

(None).

## Syslog Message

Ownership for disk %s (S/N %s) is being stolen from node "%s" (System ID: %u).

## Parameters

**diskname** (STRING): Name of the disk whose ownership is being reassigned.

**serialno** (STRING): Serial number of the disk.

**hostname** (STRING): Name of the host that currently owns the disk.

**systemid** (INT): System ID of the host that currently owns the disk.

**thread** (STRING): Name of the thread that issued the ownership change request.

**APIname** (STRING): Name of the API that issued the ownership change request.

# diskown.stealingdiskprevented events

## diskown.StealingDiskPrevented

### Severity

NOTICE

### Description

This message occurs when a "storage disk assign" operation is performed on a disk that is owned and has a reservation from another system. The reservation prevents the disk ownership change from being completed.

### Corrective Action

Perform the "storage disk assign" operation on the system that currently owns the disk.

### Syslog Message

Ownership for disk %s (S/N %s) cannot be changed because node "%s" (System ID: %u) currently owns the persistent reservation.

### Parameters

**diskname** (STRING): Name of the disk whose ownership is being assigned.

**serialno** (STRING): Serial number of the disk.

**hostname** (STRING): Name of the host that currently owns the disk.

**systemid** (INT): System ID of the host that currently owns the disk.

**thread** (STRING): Name of the thread that issued the ownership change request.

**APIname** (STRING): Name of the API that issued the ownership change request.

# diskown.uprevversion events

## diskown.UpRevVersion

### Severity

NOTICE

### Description

This message occurs when disk ownership information has been read and found to be from a version that is

higher (newer) than what the current version of Data ONTAP understands. New fields in the ownership information will be ignored.

#### **Corrective Action**

(None).

#### **Syslog Message**

(None).

#### **Parameters**

**diskname** (STRING): The name of the disk with the newer version.

**serialno** (STRING): The serial number of the disk.

**foundversion** (INT): The version of the disk ownership information found on the disk.

**curversion** (INT): The version of the disk ownership information that the current Data ONTAP uses.

## **diskown.waitingforreassign events**

### **diskown.WaitingForReassign**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a giveback is occurring and the "storage disk reassign" is in progress. The giveback must wait until the reassign operation is complete.

#### **Corrective Action**

(None).

#### **Syslog Message**

Waiting for ownership of the partner disks to be updated as a result of a "storage disk reassign" command.

#### **Parameters**

(None).

# diskshelf events

## diskshelf.psu events

### diskShelf.psu.added

#### Severity

NOTICE

#### Description

This message occurs when a power supply unit is added to the disk shelf.

#### Corrective Action

(None).

#### Syslog Message

%s power supply was added to %s.shelf%d

#### Parameters

**location** (STRING): Which power supply was added.  
**channelName** (STRING): Channel connected to disk shelf.  
**shelfIdent** (INT): ID number displayed on the disk shelf OPS panel.  
**serialNumber** (STRING): Serial number of disk shelf.

### diskShelf.psu.removed

#### Severity

ERROR

#### Description

This message occurs when a power supply unit is removed from the disk shelf.

#### Corrective Action

Insert a functional power supply unit into the disk shelf.

#### Syslog Message

%s power supply was removed from %s.shelf%d

#### Parameters

**location** (STRING): Which power supply was removed.  
**channelName** (STRING): Channel connected to disk shelf.  
**shelfIdent** (INT): ID number displayed on the disk shelf OPS panel.  
**serialNumber** (STRING): Serial number of disk shelf.

## diskshelf.support events

### diskShelf.support.file

**Severity**

NOTICE

**Description**

This message occurs when a new disk shelf support file is detected. Current support level of existing shelves will be checked.

**Corrective Action**

(None).

**Syslog Message**

New disk shelf support file (DQP) detected.

**Parameters**

(None).

## diskshelf.supported events

### diskShelf.supported

**Severity**

NOTICE

**Description**

This message occurs when a newly supported disk shelf is detected.

**Corrective Action**

(None).

**Syslog Message**

Disk shelf %s (%s) is now supported.

**Parameters**

**device** (STRING): Device address of the disk shelf.

**product\_id** (STRING): Product ID of the disk shelf.

## diskshelf.unsupported events

### diskShelf.unsupported

**Severity**

ALERT

**Description**

This message occurs when ONTAP® software detects an unsupported disk shelf.

**Corrective Action**

Verify that the disk shelf attached is supported by the version of ONTAP running on the storage controller. If the disk shelf is supported, replace the disk shelf module.

### **Syslog Message**

Disk shelf %s (%s) is unsupported.

### **Parameters**

**device** (STRING): Device address of the disk shelf.

**product\_id** (STRING): Product ID of the disk shelf.



# dlock events

## dlock.release events

### dlock.release.error

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software reports an error while attempting to release an internal lock. This might impact services like Application-aware data management.

#### Corrective Action

Run the node local "dlock show" diag-level command to confirm whether the lock still exists. If the lock exists, run the node local "dlock release" diag-level command to clear the lock. If the lock does not exist, then the error was likely transient and can be ignored. Contact NetApp technical support for more information or assistance.

#### Syslog Message

Failed to release internal lock. Lock ID: "%s". Owner ID: "%s". Reason: %s.

#### Parameters

**lockid** (STRING): Lock ID.

**ownerid** (STRING): Owner ID.

**error** (STRING): Error that was encountered.

# dm events

## dm.cluster events

### dm.cluster.rebaselinefailed

**Severity**

ERROR

**Description**

This message occurs when resynchronization of the cluster configuration failed in a MetroCluster(tm) configuration.

**Corrective Action**

Use the "metrocluster check config-cluster resync" command to resynchronize the cluster configuration.

**Syslog Message**

Cluster configuration resynchronization failed for cluster %s.

**Parameters**

**cluster** (STRING): Name of the cluster whose configuration is being replicated.

## dm.vserver events

### dm.vserver.mcc.rebaselinefailed

**Severity**

ERROR

**Description**

This message occurs when resynchronization of the Vserver configuration failed in a MetroCluster(tm) configuration.

**Corrective Action**

Use the "metrocluster vserver resync" command to resynchronize the Vserver configuration.

**Syslog Message**

Vserver configuration resynchronization failed for Vserver %s.

**Parameters**

**vserver** (STRING): Name of the Vserver whose configuration is being replicated.

### dm.vserver.sm.rebaselinefailed

**Severity**

ERROR

**Description**

This message occurs when resynchronization of the Vserver configuration failed in a Vserver DR

relationship.

**Corrective Action**

Use the "snapmirror resync" command to resynchronize the Vserver configuration.

**Syslog Message**

Vserver configuration resynchronization failed for Vserver %s.

**Parameters**

**vserver** (STRING): Name of the Vserver whose configuration is being replicated.

**dm.vserver.unlockfailed****Severity**

ERROR

**Description**

This message occurs when unlocking the Vserver failed after resynchronization of the Vserver configuration in a MetroCluster(tm) configuration or Vserver DR setup. Management operations on this Vserver will fail, until an unlock operation is done by the admin manually.

**Corrective Action**

Use the "vserver unlock -force true" command at diagnostic privilege to manually unlock the Vserver.

**Syslog Message**

Vserver unlock failed for Vserver %s.

**Parameters**

**vserver** (STRING): Name of the Vserver whose unlock operation failed.

# dml events

## dml.drive events

### dml.drive.offline.status

**Severity**

NOTICE

**Description**

This message occurs when drive offline completes.

**Corrective Action**

(None).

**Syslog Message**

Drive offline request for drive %s %s, (%d).

**Parameters**

**drive\_name** (STRING): Drive name.  
**status\_str** (STRING): Drive offline status string.  
**status\_code** (INT): Drive offline status code.

### dml.drive.online.status

**Severity**

NOTICE

**Description**

This message occurs when drive online completes.

**Corrective Action**

(None).

**Syslog Message**

Drive online request %s %s, (%d).

**Parameters**

**drive\_name** (STRING): Drive name.  
**status\_str** (STRING): Drive online status string.  
**status\_code** (INT): Drive online status code.

## dml.fae events

### dml.fae.key.import.failed

**Severity**

ERROR

## Description

This message occurs when the Device Management Layer fails to import key material.

## Corrective Action

Verify whether aggregate keys are required. Retry the key transport operation.

## Syslog Message

Key import failed. Reason: %s.

## Parameters

**reaon** (STRING): Failure reason.

# dml.init events

## dml.init.fail

### Severity

ERROR

### Description

This message occurs when the Device Management Layer (DML) fails initialization.

### Corrective Action

Retry the startup.

### Syslog Message

Firestorm initialization failed with error code (%d).

### Parameters

**error\_code** (INT): Failure code.

# dml.initialized events

## dml.initialized

### Severity

NOTICE

### Description

This message occurs when the Device Management Layer (DML) is started.

### Corrective Action

(None).

### Syslog Message

DML initialized.

### Parameters

(None).

# dms events

## dms.restserver events

### dms.restserver.ssl.disabled

**Severity**

ERROR

**Description**

This message occurs when the system hits an error in enabling SSL authentication in the rest server.

**Corrective Action**

Check that the SSL certificates and keys given to the rest server are accessible and valid.

**Syslog Message**

Failed to set SSL authentication in DMAP rest server due to %s

**Parameters**

**error** (STRING): Text describing the error.

# dns events

## dns.disabled events

### dns.disabled.configs.removed

**Severity**

NOTICE

**Description**

This message occurs when the last system to be upgraded in a cluster removes DNS configurations that were marked as disabled.

**Corrective Action**

(None).

**Syslog Message**

All DNS configurations with the "-state" parameter set to "disabled" were removed.

**Parameters**

(None).

## dns.server events

### dns.server.timed.out

**Severity**

ERROR

**Description**

This message occurs when the DNS server fails to respond to a query and timeout occurs.

**Corrective Action**

Make sure that the DNS server is up and running and that there are no networking issues preventing the Vserver from communicating with the DNS server.

**Syslog Message**

DNS server %s did not respond to vserver = %s within timeout interval.

**Parameters**

**ipaddr** (STRING): Ip address of the DNS server.

**vsName** (STRING): Name of the Vserver for which DNS response timed out.

# download events

## download.request events

### download.request

**Severity**

NOTICE

**Description**

This event is generated when a user-initiated download is initiated.

**Corrective Action**

(None).

**Syslog Message**

Operator requested download initiated

**Parameters**

(None).

## download.requestdone events

### download.requestDone

**Severity**

NOTICE

**Description**

This event is generated when a user-initiated download has completed.

**Corrective Action**

(None).

**Syslog Message**

Operator requested download completed

**Parameters**

(None).

## download.requestdoneerror events

### download.requestDoneError

**Severity**

ERROR



**Description**

This event is generated when a user-initiated download did not complete successfully because of a problem. See console log for problem and corrective action.

**Corrective Action**

(None).

**Syslog Message**

Operator requested download failed.

**Parameters**

(None).

# ds events

## ds.sas events

### ds.sas.config.warning

#### Severity

ERROR

#### Description

This message occurs when the system detects a configuration problem on the shelf I/O module.

#### Corrective Action

Reseat the disk shelf I/O module. If that does not fix the problem, replace the disk shelf I/O module.

#### Syslog Message

Configuration error on Serial Attached SCSI (SAS) I/O module %s channel %s shelf ID %lld - %s of %d expected %d

#### Parameters

**module\_id** (STRING): Module (A or B) that is experiencing the problem.  
**channel** (STRING): Channel of the shelf that is experiencing the problem.  
**shelfId** (LONGINT): Shelf ID of the shelf that is experiencing the problem.  
**desc** (STRING): Description of the parameter that is misconfigured.  
**actual** (INT): Actual value for this configuration parameter on the shelf I/O module.  
**expected** (INT): Expected value for this configuration parameter on the disk shelf I/O module.

### ds.sas.drivephy.disableErr

#### Severity

ERROR

#### Description

This message occurs when a drive physical layer device (phy) on a Serial Attached SCSI (SAS) I/O module is disabled due to one of the following reasons: manually bypassed; exceeded loss of double word synchronization threshold; exceeded running disparity threshold Transmitter fault; exceeded CRC error threshold; exceeded invalid double word threshold; exceeded phy reset problem threshold; exceeded broadcast change threshold; or mirroring disabled on the other I/O module.

#### Corrective Action

Replace the disabled disk drive.

#### Syslog Message

Disk %s%s on channels %s shelf ID %lld %s bay %s %s.

#### Parameters

**phy** (STRING): Physical layer device (phy) that is being disabled.  
**disk\_data** (STRING): Disk data includes vendor name, model, firmware revision and the serial number of the disabled disk. If unable to access the disk then, this field is blank.  
**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED

if unconnected.

**shelfID** (LONGINT): Shelf ID of the disk shelf reporting the error.

**element** (STRING): IOM (I/O module) designator (A or B).

**bay** (STRING): Bay in which the disk is installed.

**reason** (STRING): Reason the disk phy was disabled.

## **ds.sas.element.fault**

### **Severity**

ERROR

### **Description**

This message occurs when a Serial Attached SCSI (SAS) module in a disk shelf reports an error.

### **Corrective Action**

Check cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

### **Syslog Message**

State of shelf module %s is %s on channels %s shelf ID %d

### **Parameters**

**element** (STRING): Shelf module designator (A or B).

**state** (STRING): Current state of IOM.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

## **ds.sas.element.xport.error**

### **Severity**

ERROR

### **Description**

This message occurs when a Serial Attached SCSI (SAS) disk shelf module reports a transport error.

### **Corrective Action**

Check cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

### **Syslog Message**

Shelf module %s XPORT ERROR on channels %s shelf ID %d

### **Parameters**

**element** (STRING): Shelf module designator (A or B).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

## ds.sas.hostphy.disableErr

### Severity

ALERT

### Description

This message occurs when a host physical layer device (phy) on a Serial Attached SCSI (SAS) I/O module is disabled due to one of the following reasons: Invalid SAS2-SAS3 Transition; manually bypassed; exceeded loss of double word synchronization threshold; exceeded running disparity threshold; transmitter fault; exceeded CRC error threshold; exceeded invalid double word threshold; exceeded phy reset problem threshold; exceeded broadcast change threshold; or mirroring disabled on the other I/O module.

### Corrective Action

If the Host phys were disabled because of an invalid transition, remove the SAS cable on the disabled phys completely from both IOMs and then recable correctly. Removing the cable completely allows both IOMs to reset. For other disabled host phy conditions, replace the shelf module to which the affected host phy belongs.

### Syslog Message

Port %s on channels %s shelf ID %lld %s %s.

### Parameters

**phy** (STRING): Host port that is being disabled.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (LONGINT): Shelf ID of the shelf reporting the error.

**element** (STRING): I/O module designator (A or B).

**reason** (STRING): Reason the host phy was disabled.

## ds.sas.id.conflict.cleared

### Severity

NOTICE

### Description

This message occurs when a shelf ID conflict on a channel reported earlier is resolved.

### Corrective Action

(None).

### Syslog Message

Shelf ID conflict on channel %s (shelf ID: %lld, S/N: %s) was cleared.

### Parameters

**channel** (STRING): Channel that the shelf is attached to.

**shelfid** (LONGINT): Shelf ID.

**SerialNumber** (STRING): Shelf serial number.

## ds.sas.id.conflict.detected

**Severity**

ERROR

**Description**

This message occurs when a SAS shelf ID conflict is detected on a channel.

**Corrective Action**

Ensure that all of the disk shelves in a high-availability pair have a unique shelf ID. For instructions, see "Changing the disk shelf ID" in the Disk Shelf Installation and Service Guide for the shelf model.

**Syslog Message**

SAS shelf on channel %s (shelf ID: %lld, S/N: %s) reported a shelf ID conflict.

**Parameters**

**channel** (STRING): Channel that the shelf is attached to.

**shelfId** (LONGINT): Shelf ID.

**SerialNumber** (STRING): Shelf serial number.

**ds.sas.multPhys.disableErr****Severity**

ALERT

**Description**

This message occurs when the system detects that multiple Physical Layer Devices (PHYs) are disabled on a SAS disk shelf.

**Corrective Action**

Check whether the problems on the PHYs are valid by looking for PHY errors in nearby events. If multiple PHYs are disabled at the same time in a particular disk shelf module, replace the disk shelf module.

**Syslog Message**

Multiple phys have been disabled on the disk shelf %d on Channels %s, %s.

**Parameters**

**shelfID** (INT): ID of the disk shelf reporting the error.

**channels** (STRING): Channels that have the error. "Channel" can be an adapter port, (for example, 0c), or PARTNER for high-availability partner, or UNUSED if unconnected.

**element** (STRING): Disk shelf module designator (A or B).

**ds.sas.ses.disableErr****Severity**

ALERT

**Description**

This message occurs when a virtual (SES) physical layer device (phy) on a Serial Attached SCSI (SAS) I/O module is disabled due to one of the following reasons: manually bypassed; exceeded loss of double word synchronization threshold; exceeded running disparity threshold Transmitter fault; exceeded CRC error threshold; exceeded invalid double word threshold; exceeded phy reset problem threshold; or exceeded broadcast change threshold.

### Corrective Action

Replace the shelf module.

### Syslog Message

SES phy on channel %s shelf ID %lld module %s is disabled.

### Parameters

**channel** (STRING): Channel that has the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (LONGINT): Shelf ID of the shelf reporting the error.

**element** (STRING): I/O module designator (A or B).

## ds.sas.xfer.element.fault

### Severity

ERROR

### Description

This message occurs when an element had a fault during an I/O request. It might be due to a transient condition in link connectivity.

### Corrective Action

Check cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

### Syslog Message

I/O failure on shelf module %s on channels %s shelf id %d, status is %s.

### Parameters

**element** (STRING): Shelf module designator (A or B).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

**state** (STRING): Current state of IOM.

## ds.sas.xfer.not.sent

### Severity

ERROR

### Description

This message occurs when an I/O transfer could not be sent. It might be due to a transient condition in link connectivity.

### Corrective Action

Check cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

### Syslog Message

Shelf module %s I/O failure NOT\_SENT channels %s shelf ID %d

## Parameters

**element** (STRING): Shelf module designator (A or B).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

## ds.sas.xfer.unknown.error

### Severity

EMERGENCY

### Description

This message occurs when an unknown error occurs during an I/O request.

### Corrective Action

Ensure that the shelf module firmware and ONTAP® software are compatible versions. Find firmware versions at the NetApp technical support web site.

### Syslog Message

Shelf module %s I/O failure UNKNOWN: %d channels %s shelf id %d

## Parameters

**element** (STRING): Shelf module designator (A or B).

**status** (INT): Error status code.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

## ds.sas.xfer.xport.error

### Severity

ERROR

### Description

This message occurs when a transport error occurs during an I/O request. It might be due to a transient condition in link connectivity.

### Corrective Action

Check cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

### Syslog Message

Shelf module %s I/O failure TRANSPORT\_ERROR channels %s shelf ID %d.

## Parameters

**element** (STRING): Shelf module designator (A or B).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Shelf ID (0-99) of the shelf reporting the error.

# dump events

## dump.bitmap events

### dump.bitmap.fail

**Severity**

ERROR

**Description**

This event is generated during dump restarting and looking for the bitmap file and it is unable to locate

**Corrective Action**

(None).

**Syslog Message**

Missing bitmap files.

**Parameters**

(None).

## dump.fail events

### dump.fail.snapreserve

**Severity**

ERROR

**Description**

This event is generated while dump is trying to restart and there is no space left in the snap reserve

**Corrective Action**

Please increase the snap reserve space to avoid this error.

**Syslog Message**

Snapshots overflow reserved space.

**Parameters**

(None).

## dump.file events

### dump.file.has.security.label

**Severity**

NOTICE



## Description

This message occurs when the dump engine encounters an NFS security label during any type of dump operation. The dump engine does not back up any NFS security labels but skips them and continues.

## Corrective Action

(None).

## Syslog Message

The dump engine encountered one or more inodes with NFS security labels. The dump engine will back up the inodes but will not back up the associated security labels.

## Parameters

**src\_path** (STRING): Path being backed up.

# dump.inconsistent events

## dump.inconsistent.acl\_data

### Severity

ERROR

### Description

This message occurs when the dump engine encounters a disk error in an ACL file in a Snapshot(tm) copy. Backup will abort.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Dump of source path (%s), Snapshot copy ID (%d): WAFL ACL data inconsistent due to disk errors. Aborting the backup.

### Parameters

**src\_path** (STRING): Source path being backed up.

**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.

## dump.inconsistent.file\_data

### Severity

ERROR

### Description

This message occurs when the dump engine encounters a disk error in a file in a Snapshot(tm) copy. Backup will abort because the user selected the option to abort on disk error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Dump of source path (%s), Snapshot copy ID (%d), inum (%u): WAFL data inconsistent due to disk errors. Aborting the backup as directed by ndmpd.abort\_on\_disk\_error option.

## Parameters

**src\_path** (STRING): Source path being backed up.  
**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.  
**inum** (INT): Inode number of the file that encountered the disk error.

## dump.inconsistent.ind\_buf

### Severity

ERROR

### Description

This message occurs when the dump engine encounters a disk error in indirect buffers of a file in a Snapshot(tm) copy. Backup will abort because the user selected the option to abort on disk error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Dump of source path (%s), Snapshot copy ID (%d), inum (%u): WAFL metadata and data inconsistent due to disk errors. Aborting the backup as directed by ndmpd.abort\_on\_disk\_error option.

## Parameters

**src\_path** (STRING): Source path being backed up.  
**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.  
**inum** (INT): Inode number of the file that encountered the disk error.

## dump.inconsistent.stream\_data

### Severity

ERROR

### Description

This message occurs when the dump engine encounters a disk error in a stream associated with a file in a Snapshot(tm) copy. Backup will abort because the user selected the option to abort on disk error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Dump of source path (%s), Snapshot copy ID (%d), stream inum (%u), base inum(%u): WAFL data inconsistent due to disk errors. Aborting the backup as directed by ndmpd.abort\_on\_disk\_error option.

## Parameters

**src\_path** (STRING): Source path being backed up.  
**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.  
**sinum** (INT): Inode number of the stream that encountered the disk error.  
**inum** (INT): Inode number of the basefile whose associated stream encountered the disk error.

## dump.inconsistent.wafl\_metadata

**Severity**

ERROR

**Description**

This message occurs when the dump engine encounters a disk error in either the inode file, a directory or a stream directory in a Snapshot(tm) copy. As a result of this error, the backup is aborted.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Dump of source path (%s), Snapshot copy ID (%d): WAFL metadata inconsistent due to disk errors.  
Aborting the backup.

**Parameters**

**src\_path** (STRING): Source path being backed up.

**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.

## dump.inconsistent\_streamdir\_lunclone events

### dump.inconsistent\_streamdir\_lunclone

**Severity**

ERROR

**Description**

This message occurs when the dump engine encounters a disk error in a stream directory of an inconsistent LUN clone in a Snapshot(tm) copy. Backup will abort because the user selected the option to abort on disk error.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Dump of source path (%s), Snapshot copy ID (%d): WAFL metadata inconsistent due to disk errors.  
Aborting the backup.

**Parameters**

**src\_path** (STRING): Source path being backed up.

**snap\_id** (INT): ID of the Snapshot copy where the disk error was encountered.

## dump.max events

### dump.max.limit

**Severity**

ERROR

## Description

This message occurs when new dump or restore sessions cannot be started because the limit for this platform has been reached.

## Corrective Action

Dump/restore operation limit has been reached for this platform.

## Syslog Message

Maximum limit of dump/restore sessions (%d) has been reached.

## Parameters

**max\_limit** (INT): Maximum dump and restore session limits.

# dump.notify events

## dump.notify.backward.ctime

### Severity

ALERT

### Description

This message occurs when the dump incremental backup might be incomplete. When setting ctime to past values is allowed, files copied to volumes through SMB/CIFS might have timestamps that are in the past compared to the time of copy. If an inode number is reused for the files with timestamps in the past that were copied into the volume after the last backup, the files will be missed in the incremental backup.

### Corrective Action

Disable updates to ctime with past values by setting "bootarg.wafl.allow\_ctime\_backward" to "false", reboot the nodes, and then start a new backup chain from the baseline. For more information on the workaround, contact Technical Support.

### Syslog Message

The incremental backup (Dump name = "%s", Snapshot copy ID = "%d") might be incomplete because setting ctime to past values is allowed on this node.

### Parameters

**dump\_name** (STRING): Dump name.

**snap\_id** (INT): ID of the Snapshot copy being backed up.

# dump.register events

## dump.register.commit

### Severity

ERROR

### Description

This event is generated when a registry transaction failed to commit

**Corrective Action**

(None).

**Syslog Message**

Failed to flush backup table entry upgrade to registry: %s

**Parameters**

**subject** (STRING): Error description

**dump.register.create****Severity**

ERROR

**Description**

This event is generated when a registry creation is failed. Memory allocation failure

**Corrective Action**

(None).

**Syslog Message**

Failed to create a registry transaction for dump(%d).

**Parameters**

**dump\_id** (INT): Dump ID

**dump.register.set****Severity**

ERROR

**Description**

This event is generated when a registry is tried to set and failed because a freed transaction is being tried to set

**Corrective Action**

(None).

**Syslog Message**

Failed to write context update information to registry.

**Parameters**

(None).

**dump.snapshot events****dump.snapshot.fail****Severity**

ERROR

## Description

This event is generated during dump restarting and looking for the snapshot that was created by the dump while it was terminating last time and we are unable to locate the snapshot

## Corrective Action

(None).

## Syslog Message

Unable to locate snapshot needed by dump context ID (%d).

## Parameters

**dump\_id** (INT): Dump ID

# dump.terminate events

## dump.terminate

### Severity

NOTICE

### Description

This event is generated during dump termination, we see if we can retrieve the name of the dump. If we are unable to retrieve the name of the dump then dump id that is terminated is notified or else both dump id and name of the dump are notified

### Corrective Action

(None).

### Syslog Message

Terminating dump(%d)%s

### Parameters

**dump\_id** (INT): Dump ID

**dump\_name** (STRING): Dump Name

## dump.terminate.reserved

### Severity

NOTICE

### Description

This event is generated when we tried to terminate dump and have checked if dump is in the reserved state. This is just a notification message for dump termination and it is also restartable.

### Corrective Action

(None).

### Syslog Message

Terminating restartable dump(%d)

## Parameters

`dump_id` (INT): Dump ID

# dump.transit events

## dump.transit.fail

### Severity

ERROR

### Description

This event is generated when the failover is still in transit while waiting for the dump to be initialized

### Corrective Action

(None).

### Syslog Message

Failover is still in transit.

## Parameters

(None).

# dump.volume events

## dump.volume.memory

### Severity

NOTICE

### Description

This event is generated when we are trying to dump a volume which is full and we cannot write the offsetmap on to the disk and we are using memory instead of disk to do this.

### Corrective Action

Please free up some of the disk space for this volume in order to avoid this memory consumption.

### Syslog Message

DUMP using memory due to lack of disk space on /vol/%s.

## Parameters

`vol_name` (STRING): Name of the volume that is full which needs to be freed

# dump.xfercomplete events

## dump.xfercomplete.inconsistent\_data

### Severity

ERROR

**Description**

This message occurs when the dump engine completes data backup with disk errors in user files.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Dump of source path (%s): One or more files in the backup encountered disk errors. Refer to /etc/log/backup for a list of files that encountered disk errors.

**Parameters**

**src\_path** (STRING): Source path being backed up.



# ekmip events

## ekmip.akv events

### ekmip.akv.available

#### Severity

NOTICE

#### Description

This message occurs when a configured Azure Key Vault (AKV) that was previously reported unavailable for key operations is now available.

#### Corrective Action

(None).

#### Syslog Message

The AKV Key Management Service is now available. AKV configuration details, SVM: "%s", configuration name: "%s", AKV name: "%s".

#### Parameters

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**configName** (STRING): Name of the AKV configuration.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

### ekmip.akv.keyExpired

#### Severity

ERROR

#### Description

This message occurs when the current key version associated with the key identifier of the configured Azure Key Vault (AKV) has expired. The key version might be out of compliance. It is recommended that you update the key version to an active, non-expired key. ONTAP is still able to unwrap the top-level internal key protection key with the expired key version and there is no interruption to data availability.

#### Corrective Action

Update the key-ID to an active, non-expired key at the AKV portal.

#### Syslog Message

The AKV key version has expired. Key ID: "%s". AKV configuration details, SVM: "%s", configuration name: "%s", AKV name: "%s".

#### Parameters

**keyId** (STRING): Key ID.

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**configName** (STRING): Name of the AKV configuration.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

## ekmip.akv.keyVersionChanged

### Severity

NOTICE

### Description

This message occurs when the current key version associated with the key identifier of the configured Azure Key Vault (AKV), the AKV key-ID, changes. The key version might have been auto-rotated or manually changed at the AKV. Additionally, the top-level internal key protection key has been successfully re-wrapped with the current key version of the AKV key-ID.

### Corrective Action

(None).

### Syslog Message

The AKV key version has been changed. AKV configuration details, SVM "%s", configuration name "%s", AKV: "%s", previous key version: "%s", current key version: "%s".

### Parameters

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**configName** (STRING): Name of the AKV configuration.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

**oldKeyVersion** (STRING): Previous key version.

**newKeyVersion** (STRING): Current key version.

## ekmip.akv.notAvailable

### Severity

ALERT

### Description

This message occurs when a check for the availability of the configured Azure Key Vault (AKV) for key operations fails. The AKV might be down or there might be a network-related problem preventing communication. Without access to the AKV, the node might not be able to restore authentication keys needed to unlock NetApp Storage Encryption (NSE) or SED drives, or encryption keys needed to mount encrypted volumes. If the AKV is unavailable the next time this node boots, then the failure to restore the keys might prevent the node from booting successfully or prevent the encrypted volumes hosted on this node from coming online.

### Corrective Action

Ensure that the AKV configuration is correct by using the "security key-manager external azure show -vserver <vserver>" command and verifying that the cluster state is available. If it is not available, run the "security key-manager external azure restore" command to bring the cluster state to available. If the issue still persists, contact technical support.

### Syslog Message

The AKV Key Management Service is not available. AKV configuration details, SVM: "%s", configuration name: "%s", AKV name: "%s".

### Parameters

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**configName** (STRING): Name of the AKV configuration.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

## ekmip.akv.volOffline

### Severity

ALERT

### Description

This message occurs when the SVM has been blocked and any encrypted volumes belonging to the SVM have been taken offline due to errors received when attempting to access the key owned by the Azure Key Vault (AKV). Reasons for AKV key access errors include a disabled key, a key not being found and a key missing encryption and decryption privileges. ONTAP polls the AKV every 15 minutes to verify key accessibility. If, for 60 minutes, ONTAP does not receive a successful response to a poll, all encrypted volumes are taken offline and remain offline until the key access issues are resolved at the AKV. When ONTAP does receive a successful response, the volumes are brought back online automatically.

### Corrective Action

Resolve the key access issues at the AKV portal. Ensure that the key is active and has the required encryption and decryption privileges.

### Syslog Message

The SVM has been blocked and any encrypted volumes belonging to the SVM have been taken offline due to key access errors. AKV configuration details, SVM: "%s", configuration name "%s", key ID: "%s", AKV name: "%s".

### Parameters

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**keyId** (STRING): Key ID.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

**configName** (STRING): Name of the AKV configuration.

## ekmip.akv.volOnline

### Severity

NOTICE

### Description

This message occurs when the volumes configured on the Azure Key Vault (AKV) that were unavailable due to key access issues are now back online.

### Corrective Action

(None).

### Syslog Message

Encrypted volumes belonging to the SVM are now back online. AKV configuration details, SVM: "%s", configuration name: "%s", key ID "%s", AKV name: "%s".

### Parameters

**svm** (STRING): Name of the storage virtual machine (SVM) where the AKV is configured.

**name** (STRING): Uniform resource identifier (URI) of the AKV.

**configName** (STRING): Name of the AKV configuration.

**keyId** (STRING): Key ID.

# ekmip.awskms events

## ekmip.awskms.available

### Severity

NOTICE

### Description

This message occurs when a configured Amazon Web Services (AWS) Key Management Service (KMS) that was unavailable for key operations is now available.

### Corrective Action

(None).

### Syslog Message

The AWS KMS with region: "%s" and key ID: "%s" configured for Vserver "%s" is now available.

### Parameters

**region** (STRING): AWS KMS region.

**keyId** (STRING): Key ID.

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

## ekmip.awskms.keyChg

### Severity

NOTICE

### Description

This message occurs when the key material associated with the key ID of the configured Amazon Web Services (AWS) Key Management Service (KMS) changes. The key material might have been auto-rotated or manually changed at the AWS KMS. Additionally, as a result of updated key material, the top-level internal key-protection key is successfully rewrapped with the current key material of the AWS KMS key ID.

### Corrective Action

(None).

### Syslog Message

The key material for key "%s" owned by AWS KMS region "%s" and configured for Vserver "%s" has changed.

### Parameters

**keyId** (STRING): Key ID.

**region** (STRING): AWS KMS region.

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

## ekmip.awskms.keyVersionChg

### Severity

NOTICE

## Description

This message occurs when the key version associated with the key ID of the configured Amazon Web Services (AWS) Key Management Service (KMS) changes. The key version might have been auto-rotated or manually changed at the AWS KMS. Additionally, as a result of updated key version, the top-level internal key-protection key is successfully rewrapped with the current key version of the AWS KMS key ID.

## Corrective Action

(None).

## Syslog Message

The version of the key "%s" owned by AWS KMS region "%s", configured for Vserver "%s", is changed from "%s" to "%s".

## Parameters

**keyId** (STRING): Key ID.

**region** (STRING): AWS KMS region.

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

**oldKeyVersion** (STRING): Prior key version.

**newKeyVersion** (STRING): New key version.

## ekmip.awskms.notAvailable

### Severity

ALERT

## Description

This message occurs when a check for the availability of the configured Amazon Web Services (AWS) Key Management Service (KMS) fails. The AWS KMS might be down or there might be a network-related problem preventing communication. Without access to the AWS KMS, the node might be unable to restore authentication keys needed to unlock NetApp Storage Encryption (NSE) drives or encryption keys needed to mount encrypted volumes. If the AWS KMS is unavailable, then the next time this node boots, the failure to restore the keys might prevent the node from booting successfully or the encrypted volumes hosted on this node from coming online.

## Corrective Action

Verify that the AWS KMS configuration is correct by using the "security key-manager external aws show -vserver <vserver\_name>" command, and verify that the AWS KMS cluster state is available. If the cluster state is unavailable, run the "security key-manager external aws restore" command to bring the cluster state to available. If the issue still persists, contact NetApp technical support.

## Syslog Message

The AWS KMS with region: %s and key ID: %s configured for Vserver "%s" is not available.

## Parameters

**region** (STRING): AWS KMS region.

**keyId** (STRING): Key ID.

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

## ekmip.awskms.volOffline

## Severity

ALERT

## Description

This message occurs when the Vserver has been blocked and any encrypted volumes belonging to the Vserver have been taken offline due to errors received when trying to access the key owned by the Amazon Web Services (AWS) Key Management Service (KMS). Reasons for AWS KMS key access errors include a disabled key, a destroyed key, key not found, or a user not having sufficient privileges to access the key. ONTAP polls the AWS KMS to verify key accessibility. If, after 60 minutes, ONTAP has not received a successful response to a poll, volumes are taken offline and remain offline until the key access issues are resolved at the AWS KMS. Subsequently, if ONTAP does receive a successful response, the volumes will be brought back online automatically.

## Corrective Action

Resolve the key access issues at the AWS KMS portal. Ensure that the key is active and the user has the required privileges to access the key.

## Syslog Message

The Vserver has been blocked and any encrypted volumes belonging to Vserver "%s" have been taken offline due to key access errors associated with key "%s" owned by AWS KMS region "%s".

## Parameters

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

**keyId** (STRING): Key ID.

**region** (STRING): AWS KMS region.

## ekmip.awskms.volOnline

### Severity

NOTICE

### Description

This message occurs when the volumes configured on the Amazon Web Services (AWS) Key Management Service (KMS) that were unavailable due to key access issues are now back online.

### Corrective Action

(None).

### Syslog Message

Encrypted volumes belonging to Vserver "%s" associated with key-id "%s" owned by Amazon Web Services (AWS) KMS region "%s" are now back online.

### Parameters

**vserver** (STRING): Name of the Vserver where the AWS KMS is configured.

**keyId** (STRING): Key ID.

**region** (STRING): AWS KMS region.

## ekmip.gckms events

## ekmip.gckms.available

### Severity

NOTICE

### Description

This message occurs when a configured Google Cloud Key Management Service (GCKMS) that was previously reported as unavailable for key operations is now available.

### Corrective Action

(None).

### Syslog Message

The GCKMS with project ID: %s, key ring location: %s, key ring: %s and key name: %s configured for Vserver "%s" is now available.

### Parameters

**projectId** (STRING): Project ID.

**keyringLocation** (STRING): Location of the key ring.

**keyringName** (STRING): Name of the key ring.

**keyName** (STRING): Name of the key.

**vserver** (STRING): Name of the Vserver where the GCKMS is configured.

## ekmip.gckms.keyVersionChg

### Severity

NOTICE

### Description

This message occurs when a change is made to the current key version associated with the key name of the configured Google Cloud KMS (GCKMS). The key version might have been auto-rotated or manually changed at the GCKMS. Additionally, the top-level internal key-protection key has been successfully re-wrapped with the current key version of the GCKMS key name.

### Corrective Action

(None).

### Syslog Message

The version of the key "%s" owned by Google Cloud KMS Project "%s", with key ring "%s" at "%s", configured for Vserver "%s", has been changed from "%s" to "%s".

### Parameters

**keyName** (STRING): Name of the key.

**projectId** (STRING): Project ID of the GCKMS.

**keyringName** (STRING): Name of the key ring.

**keyringLocation** (STRING): Location of the key ring.

**vserver** (STRING): Name of the Vserver where the GCKMS is configured.

**oldKeyVersion** (STRING): Prior key version.

**newKeyVersion** (STRING): New key version.

## ekmip.gckms.notAvailable

### Severity

ALERT

### Description

This message occurs when a check for the availability of the configured Google Cloud Key Management Service (GCKMS) fails. The GCKMS might be down, there might be a network related problem preventing communication, or there might be an error in the GCKMS key manager configuration. Without access to the GCKMS, one or more nodes might not be able to restore the authentication keys needed to unlock NSE drives or the encryption keys needed to mount encrypted volumes. If the GCKMS is not available to a node, then the next time the node boots, the failure to restore the keys might prevent the node from booting successfully or the encrypted volumes hosted on the node from coming online.

### Corrective Action

Verify that the GCKMS configuration is correct by using the "security key-manager external gcp check -vserver <vserver\_name>" command. If the issue still persists, contact technical support.

### Syslog Message

The GCKMS with project ID: %s, key ring location: %s, key ring: %s and key name: %s configured for Vserver "%s" is not available.

### Parameters

**projectId** (STRING): Project ID.

**keyringLocation** (STRING): Location of the key ring.

**keyringName** (STRING): Name of the key ring.

**keyName** (STRING): Name of the key.

**vserver** (STRING): Name of the Vserver where the GCKMS is configured.

## ekmip.gckms.volOffline

### Severity

ALERT

### Description

This message occurs when the Vserver has been blocked and any encrypted volumes belonging to the Vserver have been taken offline due to errors received when trying to access the key owned by the Google Cloud KMS (GCKMS). Reasons for GCKMS key access errors include a disabled key, a destroyed key, key not found, or a user not having sufficient privileges to access the key. ONTAP polls the GCKMS every 15 minutes to verify key accessibility. If, after 60 minutes, ONTAP has not received a successful response to a poll, volumes are taken offline and remain offline until the key access issues are resolved at the GCKMS. Subsequently, if ONTAP does receive a successful response, the volumes will be brought back online automatically.

### Corrective Action

Resolve the key access issues at the GCKMS portal. Ensure that the key is active and the user has the required privileges to access the key.

### Syslog Message

The Vserver has been blocked and any encrypted volumes belonging to Vserver "%s" have been taken offline due to key access errors associated with key "%s" owned by Google Cloud KMS Project "%s" with key ring location "%s" and key ring "%s".



## Parameters

**vserver** (STRING): Name of the Vserver where the GCKMS is configured.

**keyName** (STRING): Name of the key.

**projectId** (STRING): ID of the GCKMS project.

**keyringLocation** (STRING): Location of the key ring.

**keyringName** (STRING): Name of the key ring.

## ekmip.gckms.volOnline

### Severity

NOTICE

### Description

This message occurs when the volumes configured on the Google Cloud KMS (GCKMS) that were unavailable due to key access issues are now back online.

### Corrective Action

(None).

### Syslog Message

Encrypted volumes belonging to Vserver "%s" associated with key-id "%s" owned by Google Cloud KMS Project "%s" with key ring location "%s" and key ring "%s" are now back online.

## Parameters

**vserver** (STRING): Name of the Vserver where the GCKMS is configured.

**keyName** (STRING): Name of the key.

**projectId** (STRING): ID of the GCKMS project.

**keyringLocation** (STRING): Location of the key ring.

**keyringName** (STRING): Name of the key ring.

## ekmip.ikpkms events

### ekmip.ikpkms.200.missingField

### Severity

ALERT

### Description

This message occurs when trying to access the key owned by the IBM Key Protect (IKP) Key Management Service (KMS) and an "OK" response (code 200) is received but one field of the response is not populated. Reasons for IKP KMS key access errors include a disabled key, a destroyed key, key not found, or a user not having sufficient privileges to access the key. ONTAP polls the IKP KMS to verify key accessibility. The missing field is integral in determining the status of the key.

### Corrective Action

The IKP KMS provided a response with a status of 200 (OK) that did not contain all of the expected fields. Contact technical support to resolve this issue.

### Syslog Message

A 200 OK response was received from IKP KMS but the field "%s" was missing for Vserver "%s" with Keymanager Name "%s" associated with key "%s" owned by IKP KMS region "%s" with instance ID "%s".

## Parameters

**missingField** (STRING): Name of the field that is not populated.  
**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): The name of the IKP KMS.  
**keyId** (STRING): The Key ID associated with this IKP KMS instance.  
**region** (STRING): The region associated with this IKP KMS instance.  
**instanceId** (STRING): IKP KMS Instance ID.

## ekmip.ikpkms.available

### Severity

NOTICE

### Description

This message occurs when a configured IBM Key Protect (IKP) Key Management Service (KMS) that was unavailable for key operations is now available.

### Corrective Action

(None).

### Syslog Message

The IKP KMS with region: "%s", instance ID: "%s", and key ID: "%s" configured for Vserver "%s" with Keymanager Name "%s" is now available.

## Parameters

**region** (STRING): The region associated with this IKP KMS instance.  
**instanceId** (STRING): IKP KMS Instance ID.  
**keyId** (STRING): The Key ID associated with this IKP KMS instance.  
**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): The name of the IKP KMS.

## ekmip.ikpkms.keyVersionChg

### Severity

NOTICE

### Description

This message occurs when the key version associated with the key ID of the configured IBM Key Protect (IKP) Key Management Service (KMS) changes. The key version might have been auto-rotated or manually changed at the IKP KMS. Additionally, as a result of updated key version, the top-level internal key-protection key has been successfully rewrapped with the current key version of the IKP KMS key ID.

### Corrective Action

(None).

### Syslog Message

The version of the key "%s" owned by IKP KMS region "%s" and instance ID "%s", configured for Vserver "%s" with Keymanager Name "%s", has been changed from "%s" to "%s".

## Parameters

**keyId** (STRING): The Key ID associated with this IKP KMS instance.  
**region** (STRING): The region associated with this IKP KMS instance.

**instanceId** (STRING): IKP KMS Instance ID.  
**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): The name of the IKP KMS.  
**oldKeyVersion** (STRING): Prior key version.  
**newKeyVersion** (STRING): New key version.

## ekmip.ikpkms.notAvailable

### Severity

ALERT

### Description

This message occurs when a check for the availability of the configured IBM Key Protect (IKP) Key Management Service (KMS) fails. The IKP KMS might be down or there might be a network-related problem preventing communication. Without access to the IKP KMS, the node might be unable to restore authentication keys needed to unlock NetApp Storage Encryption (NSE) drives or encryption keys needed to mount encrypted volumes. If the IKP KMS is unavailable, then the next time this node boots, the failure to restore the keys might prevent the node from booting successfully or the encrypted volumes hosted on this node from coming online.

### Corrective Action

Verify that the IKP KMS configuration is correct by using the "security key-manager external ikp show -vserver <vserver\_name>" command, and verify that the IKP KMS cluster state is available. If the cluster state is unavailable, run the "security key-manager external ikp restore" command to bring the cluster state to available. If the issue still persists, contact NetApp technical support.

### Syslog Message

The IKP KMS with region: "%s", instance ID: "%s", and key ID: "%s" configured for Vserver "%s" with Keymanager Name "%s" is not available.

### Parameters

**region** (STRING): The region associated with this IKP KMS instance.  
**instanceId** (STRING): IKP KMS Instance ID.  
**keyId** (STRING): The Key ID associated with this IKP KMS instance.  
**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): The name of the IKP KMS.

## ekmip.ikpkms.volOffline

### Severity

ALERT

### Description

This message occurs when the volumes belonging to the Vserver are taken offline due to errors received when trying to access the key owned by the IBM Key Protect (IKP) Key Management Service (KMS). Reasons for IKP KMS key access errors include a disabled key, a destroyed key, key not found, or a user not having sufficient privileges to access the key. ONTAP polls the IKP KMS to verify key accessibility. If, after 60 minutes, ONTAP has not received a successful response to a poll, volumes are taken offline and remain offline until the key access issues are resolved at the IKP KMS. Subsequently, if ONTAP does receive a successful response, the volumes will be brought back online automatically.

### Corrective Action

Resolve the key access issues at the IKP KMS portal. Ensure that the key is active and the user has the required privileges to access the key.

### Syslog Message

Encrypted volumes belonging to Vserver "%s" with Keymanager Name "%s" were taken offline due to key access errors associated with key "%s" owned by IKP KMS region "%s" and instance ID "%s".

### Parameters

**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): The name of the IKP KMS.  
**keyId** (STRING): The Key ID associated with this IKP KMS instance.  
**region** (STRING): The region associated with this IKP KMS instance.  
**instanceId** (STRING): IKP KMS Instance ID.

## ekmip.ikpkms.volOnline

### Severity

NOTICE

### Description

This message occurs when the volumes configured on the IBM Key Protect (IKP) Key Management Service (KMS) that were unavailable due to key access issues are now back online.

### Corrective Action

(None).

### Syslog Message

Encrypted volumes belonging to Vserver "%s" with keymanager name "%s" associated with key-id "%s" owned by IBM Key Protect (IKP) KMS region "%s" and instance ID "%s" are now back online.

### Parameters

**vserver** (STRING): Name of the Vserver where the IKP KMS is configured.  
**kmName** (STRING): Name of the IKP KMS.  
**keyId** (STRING): Key ID associated with this IKP KMS instance.  
**region** (STRING): Region associated with this IKP KMS instance.  
**instanceId** (STRING): IKP KMS Instance ID.

# ems events

## ems.asup events

### ems.asup.noteto.transition

#### Severity

NOTICE

#### Description

This message occurs when the e-mail addresses for deprecated AutoSupport "noteto" messages are copied to EMS event notification destinations during an upgrade to Data ONTAP® 9.0 or later. This results in duplicate event notifications: one EMS notification and one AutoSupport short note message. Use the "event notification destination {show|modify}" command to view or modify deprecated "noteto" e-mail addresses to use new e-mail destinations. This copy action also creates an event notification that maps the "important-events" filter to all of the newly created destinations. Use the "event notification {show|modify}" command to view or modify the new event notification, as necessary. All events that match the "important-events" filter send notifications to the new destinations, including the AutoSupport triggers with "noteto" enabled.

#### Corrective Action

To remove event message duplication, remove the AutoSupport "noteto" addresses by entering the "system node autosupport modify -node \* -noteto -" command.

#### Syslog Message

AutoSupport -noteto e-mail addresses have been copied and transitioned to EMS event notification destinations.

#### Parameters

(None).

## ems.engine events

### ems.engine.event.dnsLkupFail

#### Severity

ERROR

#### Description

This message occurs when the Event Management System (EMS) fails to send an event to a mail destination due to name resolution lookup failures for the mail server specified in EMS configuration options. Possible reasons for failure: DNS servers are not reachable, name services might not be configured correctly or, the configured mail server host name is no longer valid.

#### Corrective Action

Use the "vserver services name-service dns show" command to check the DNS configuration. If DNS is not configured correctly, use the "vserver services name-service dns modify" command to correct DNS server information, or the "vserver services name-service dns host modify" to correct static host name-to-IP address mapping information. Use the "event config {show|modify}" command to verify or change the configured mail server.

## Syslog Message

Failed to send EMS event %s to mail server "%s" because the host name could not be resolved.

## Parameters

**emsId** (STRING): EMS identifier of the event that was not sent to the mail destination.

**mailserver** (STRING): Mail server specified in EMS configuration options that could not be resolved.

## ems.engine.event.undefinedEvent

### Severity

NOTICE

### Description

This message occurs when the Event Management System (EMS) receives an event that it does not recognize. The event cannot be processed.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**requestedEventName** (STRING): EMS identifier of the event that was undefined.

**requestedEventParams** (STRING): List of parameters provided with the requested event.

## ems.engine.event.unrecognizedEventVersion

### Severity

NOTICE

### Description

This message occurs when the Event Management System (EMS) receives an event name that it recognizes, but the event has unexpected parameters. This might occur naturally when the EMS tries to replay an event that occurred before an upgrade (or revert) and the event definition changed between versions.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**originalEventXml** (STRING): Event name and parameters that were not recognized.

## ems.engine.flush

### Severity

INFORMATIONAL

## Description

This message occurs when the event queues are flushed. Event queues are flushed on system shutdown and before some types of AutoSupport™ data collection.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

(None).

# ems.snmp events

## ems.snmp.parse.error

### Severity

ERROR

### Description

This message occurs when SNMP parsing of an input packet fails. It could fail for various reasons. The most common failures are due to invalid values for the integer and string fields in the Protocol Data Unit (PDU) and message overflows.

### Corrective Action

An agent is incorrectly communicating with the system using SNMP. Find and disable or correct this agent. An Ethernet "sniffer" might be required to identify the client source that is sending the unsupported or incorrect SNMP request.

### Syslog Message

%s

### Parameters

**errorMsg** (STRING): Returned error message from the parsing of the input packet.

# ems.test events

## ems.test.notice

### Severity

NOTICE

### Description

This message occurs when the system generates a test EMS message of notice severity in response to a user test request.

### Corrective Action

(None).

**Syslog Message**

(None).

**Parameters**

(None).



# enc events

## enc.operation events

### enc.operation.aborted

#### Severity

ERROR

#### Description

This message occurs when either the rekeying of an encrypted flexible volume, or the conversion of the volume from unencrypted to encrypted, has been aborted.

#### Corrective Action

Determine the cause of the error by using either the "volume encryption rekey show -vserver <vserver-name> -volume <volume-name>" command or the "volume encryption conversion show -vserver <vserver-name> -volume <volume-name>" command. Correct the error. If the error was related to rekeying, continue the operation by using the "volume encryption rekey resume -vserver <vserver-name> -volume <volume-name>" command. If the error was related to conversion to encryption, continue the operation by using the "volume encryption conversion resume -vserver <vserver-name> -volume <volume-name>" command.

#### Syslog Message

%s scan is aborted in the volume '%s' (aggr '%s'). Reason: %s.

#### Parameters

**op** (STRING): Rekey or conversion operation.  
**volume** (STRING): Volume name.  
**aggr** (STRING): Name of the containing aggregate.  
**status** (STRING): Status of the scan.

### enc.operation.completed

#### Severity

NOTICE

#### Description

This message occurs when rekey or conversion to encryption operation is completed in a flexible volume.

#### Corrective Action

(None).

#### Syslog Message

%s operation is completed in the volume '%s' (aggr '%s').

#### Parameters

**op** (STRING): Rekey or conversion operation.  
**volume** (STRING): Volume name.  
**aggr** (STRING): Name of the containing aggregate.

## enc.operation.disableAid

### Severity

NOTICE

### Description

This message occurs when the system automatically disables the Aggregate Inline Deduplication (AID) feature for a volume as it is converted from a non-encrypted volume to an encrypted volume because AID is not supported for an encrypted volume.

### Corrective Action

(None).

### Syslog Message

AID is disabled for volume %s within aggregate %s during conversion to encryption operation.

### Parameters

**volume** (STRING): Volume name.

**aggr** (STRING): Name of the containing aggregate.

## enc.operation.failure

### Severity

ERROR

### Description

This message occurs when there is a failure in encryption operation. The operation description explains the nature of operation and reason for failure.

### Corrective Action

Ensure that there is enough space in the aggregate, cluster is healthy, volume location database is online, and then retry the volume encryption operation.

### Syslog Message

Failure in %s operation on volume '%s' (aggr '%s'). Reason '%s'.

### Parameters

**op** (STRING): Rekey or conversion operation.

**volume** (STRING): Volume name.

**aggr** (STRING): Name of the containing aggregate.

**desc** (STRING): Describes the reason for the failure of operation with arguments.

## enc.operation.started

### Severity

NOTICE

### Description

This message occurs when rekey or conversion to encryption operation is started in a flexible volume.

**Corrective Action**

(None).

**Syslog Message**

%s operation is started in the volume '%s' (aggr '%s').

**Parameters**

**op** (STRING): Rekey or conversion operation.

**volume** (STRING): Volume name.

**aggr** (STRING): Name of the containing aggregate.

# encryption events

## encryption.op events

### encryption.op.completed

#### Severity

NOTICE

#### Description

This message occurs when a rekey or conversion to encryption operation is completed on a volume.

#### Corrective Action

(None).

#### Syslog Message

%s operation is completed on the volume '%s' in Vserver '%s'.

#### Parameters

**op** (STRING): Rekey or conversion operation.  
**volumename** (STRING): Name of the volume.  
**vservername** (STRING): Name of the Vserver.

### encryption.op.paused.by.user

#### Severity

NOTICE

#### Description

This message occurs when a rekey or conversion to encryption operation is paused as a result of an operator request.

#### Corrective Action

(None).

#### Syslog Message

%s operation is paused by user on the volume '%s' in Vserver '%s'.

#### Parameters

**op** (STRING): Rekey or conversion operation.  
**volumename** (STRING): Name of the volume.  
**vservername** (STRING): Name of the Vserver.

### encryption.op.paused.on.err

#### Severity

ERROR

## Description

This message occurs when a rekey or conversion to encryption operation is paused as a result of an error. If it is a rekey operation, volume will use two encryption keys until the rekey operation is completed. New writes and rekeyed blocks will use the new key, while the data written before the rekey start will still use the old key in the read path. If it is a conversion to encryption operation, new writes and converted blocks will use the key, while data written before conversion will not use any key in the read path. Data access continues to be possible, although performance might be degraded.

## Corrective Action

Determine the cause of the error by using either the "volume encryption rekey show -vserver <vserver-name> -volume <volume-name>" command or the "volume encryption conversion show -vserver <vserver-name> -volume <volume-name>" command. Correct the error. If the error was related to rekeying, continue the operation by using the "volume encryption rekey resume -vserver <vserver-name> -volume <volume-name>" command. If the error was related to conversion to encryption, continue by using the "volume encryption conversion resume -vserver <vserver-name> -volume <volume-name>" command.

## Syslog Message

%s operation on volume '%s' in Vserver '%s' is paused due to an error. Reason '%s'.

## Parameters

**op** (STRING): Rekey or conversion operation.  
**volumename** (STRING): Name of the volume.  
**vservername** (STRING): Name of the Vserver.  
**reason** (STRING): Reason for the failure.

## encryption.op.queued

### Severity

NOTICE

### Description

This message occurs when a plaintext volume is queued for encryption.

### Corrective Action

(None).

### Syslog Message

Volume %s in Vserver %s is queued for %s.

### Parameters

**volumename** (STRING): Name of the volume.  
**vservername** (STRING): Name of the Vserver.  
**op** (STRING): Conversion operation.

## encryption.op.started

### Severity

NOTICE

### Description

This message occurs when a rekey or conversion to encryption operation is started on a volume.

**Corrective Action**

(None).

**Syslog Message**

%s operation is started on the volume '%s' in Vserver '%s'.

**Parameters**

**op** (STRING): Rekey or conversion operation.

**volumename** (STRING): Name of the volume.

**vservername** (STRING): Name of the Vserver.

## encryption.queue events

**encryption.queue.op.failed****Severity**

ERROR

**Description**

This message occurs when an operation on the volume in the encryption queue failed.

**Corrective Action**

Check the reason for the failure included in the message. Correct the issue and try the operation again.

**Syslog Message**

%s operation failed on the queued volume '%s' in Vserver '%s'. Reason: %s.

**Parameters**

**op** (STRING): Conversion operation.

**volumename** (STRING): Name of the volume.

**vservername** (STRING): Name of the Vserver.

**reason** (STRING): Reason for the failure.

# esh events

## esh.async events

### esh.async.error

#### Severity

ERROR

#### Description

This message occurs when there is an error allocating memory for an asynchronous message during disk shelf discovery processing.

#### Corrective Action

If the error occurs upon addition of a disk shelf to a running system, disconnect the shelf and reattach it. If the error occurs during boot, reboot the controller.

#### Syslog Message

Could not allocate memory for an asynchronous message.

#### Parameters

(None).

## esh.auto events

### esh.auto.term.error

#### Severity

ERROR

#### Description

This message occurs when there is an error in autotermination. Major causes of this include: 1) Last disk shelf in the loop has the termination switch set improperly. 2) Communication is lost to some number of disk shelves in the loop due to an incorrect speed setting for DS14mk2 disk shelves, a bad cable connection, or unavailable power.

#### Corrective Action

Use the 'fcstat device\_map' command to verify that all disk shelves in the loop are still accessible. If the disk shelf reporting the autotermination event is not the last disk shelf in the loop, a problem exists with some number of disk shelves in the loop. Check the speed, power, and cabling of all disk shelves in the loop after the disk shelf reporting the autotermination.

#### Syslog Message

AUTO TERMINATION ERROR: TERM switch set improperly on ESH %s on channels %s disk shelf ID %d.

#### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

# esh.bypass events

## esh.bypass.err.disk

### Severity

ERROR

### Description

This message occurs when there is a disk bypass event. It indicates that a disk is bypassed by the embedded switched hub (ESH) due to one of the following reasons: Manually bypassed; Transmitter fault; LIP F8; Data timeout; Receiver loss of signal; Comma loss of signal; Loop test before insert failure; Redundant port configuration; Exceeded stall threshold; Exceeded word error threshold; Exceeded CRC error threshold; or Exceeded clock delta error threshold.

### Corrective Action

Replace the bypassed disk.

### Syslog Message

Disk %s%son channels %s disk shelf ID %d %s%s%s.

### Parameters

**diskID** (STRING): Disk that is being bypassed.

**disk\_data** (STRING): disk data includes vendor name, model, firmware revision and the serial number of the disabled disk. If won't be able to access the disk then this field will be blank.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**bay** (STRING): Bay in which the disk is installed.

**reason** (STRING): Reason the disk was bypassed.

## esh.bypass.host.port

### Severity

ALERT

### Description

This message occurs when there is a host port bypass event. It indicates that a host port is bypassed by the embedded switched hub (ESH) due to one of the following reasons: Manually bypassed; Transmitter fault; LIP F8; Data timeout; Receiver loss of signal; Comma loss of signal; Loop test before insert failure; Redundant port configuration; Exceeded stall threshold; Exceeded word error threshold; Exceeded CRC error threshold; or Exceeded clock delta error threshold.

### Corrective Action

Replace the disk shelf module to which the concerned host port belongs.

### Syslog Message

%s on channels %s disk shelf ID %d %s%s.

### Parameters

**portID** (STRING): Host port that is being bypassed.



**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**reason** (STRING): Reason the host port was bypassed.

## esh.element events

### esh.element.fault

#### Severity

ERROR

#### Description

This message occurs when the filer cannot communicate with the ESH module.

#### Corrective Action

Check the cabling to the disk shelf. Verify that the SCSI Enclosure Services (SES) disk drives (0-1) are properly inserted. Check the status LED on the disk shelf, making sure that the fault LEDs are not lit. Clear any fault condition, if possible. Refer to the quick reference card beneath the disk shelf for information about the meanings of the LEDs.

#### Syslog Message

State of ESH module %s is %s on channels %s disk shelf id %d

#### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**state** (STRING): Current state of ESH Module.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

### esh.element.xport.error

#### Severity

ERROR

#### Description

This message occurs when there is a transport error.

#### Corrective Action

Check the cabling to the disk shelf. Verify that the SCSI Enclosure Services (SES) disk drives (0-1) are properly inserted. Check the status LED on the disk shelf, making sure that the fault LEDs are not lit. Clear any fault condition, if possible. Refer to the quick reference card beneath the disk shelf for information about the meanings of the LEDs.

#### Syslog Message

ESH %s XPORT ERROR on channels %s disk shelf ID %d.

#### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

## esh.init events

### esh.init.error

#### Severity

ERROR

#### Description

This message occurs when the system cannot allocate memory for a required service. The controller might not have sufficient memory to run the operating system, or installed memory is not functioning properly.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Could not allocate memory for ESH globals.

#### Parameters

(None).

## esh.module events

### esh.module.multiprob

#### Severity

ERROR

#### Description

This message occurs when the ESH might have failed because it is reporting multiple problems.

#### Corrective Action

Replace the ESH module if the reported multiple problems are not valid.

#### Syslog Message

channels %s disk shelf ID %d %s%s.

#### Parameters

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**reason** (STRING): Reason the message is being sent.

## esh.pctl events

## esh.pctl.bpCPLDFailed

### Severity

ERROR

### Description

This message occurs when a 'power control' command cannot be performed on the specified disk device because the backplane CPLD (Complex Programmable Logic Device) is not operational.

### Corrective Action

Power-cycle the disk shelf. If the problem persists, replace the disk shelf.

### Syslog Message

Cannot execute a 'power control' command on disk %d [%s:%s:%s] serial number [%s] in disk shelf %d on channel %s. Backplane CPLD (Complex Programmable Logic Device) in the disk shelf is not operational.

### Parameters

**loopID** (INT): Loop ID of the disk for which the 'power control' command failed.

**vendor** (STRING): Vendor name of the disk.

**model** (STRING): Model name of the disk.

**fwrev** (STRING): Firmware revision of the disk.

**serno** (STRING): Serial number of the disk.

**shelfID** (INT): Disk shelf ID of the disk shelf that reported an error.

**channelName** (STRING): Name of the Fibre Channel adapter port.

## esh.pctl.connFailed

### Severity

ERROR

### Description

This message occurs when a 'power control' command cannot be performed on the specified disk device because the CPLD (Complex Programmable Logic Device) serial link on the specified module has a problem.

### Corrective Action

Replace the specified I/O module.

### Syslog Message

Cannot execute a 'power control' command on disk %d [%s:%s:%s] serial number [%s] in disk shelf %d on channel %s. The CPLD (Complex Programmable Logic Device) serial link on I/O module %s has a problem.

### Parameters

**loopID** (INT): Loop ID of the disk for which the 'power control' command failed.

**vendor** (STRING): Vendor name of the disk.

**model** (STRING): Model name of the disk.

**fwrev** (STRING): Firmware revision of the disk.

**serno** (STRING): Serial number of the disk.

**shelfID** (INT): Disk shelf ID of the disk shelf that reported an error.

**channelName** (STRING): Name of the Fibre Channel adapter port.

**module** (STRING): Name of the module with a bad CPLD serial link.

## esh.pctl.modCPLDFailed

### Severity

ERROR

### Description

This message occurs when a 'power control' command cannot be performed on the specified disk device because the CPLD (Complex Programmable Logic Device) on an I/O module is not operational.

### Corrective Action

Replace the specified I/O module.

### Syslog Message

Cannot execute a 'power control' command on disk %d [%s:%s:%s] serial number [%s] in disk shelf %d on channel %s. The CPLD (Complex Programmable Logic Device) on I/O module %s is not operational.

### Parameters

**loopID** (INT): Loop ID of the disk for which the 'power control' command failed.

**vendor** (STRING): Vendor name of the disk.

**model** (STRING): Model name of the disk.

**fwrev** (STRING): Firmware revision of the disk.

**serno** (STRING): Serial number of the disk.

**shelfID** (INT): Disk shelf ID of the disk shelf that reported an error.

**channelName** (STRING): Name of the Fibre Channel adapter port.

**module** (STRING): I/O module with the failed CPLD.

## esh.speed events

### esh.speed.error

#### Severity

ERROR

#### Description

This message occurs when there is an error in the link speed of the disk shelves that are attached to the node. Major causes of this include: 1) Last disk shelf in the loop has the termination switch set improperly. 2) Communication is lost to some number of disk shelves in the loop due to an incorrect speed setting for DS14mk2 disk shelves, a bad cable connection, or unavailable power.

#### Corrective Action

Use the 'fcstat device\_map' command to verify that all disk shelves in the loop are still accessible. Check the speed, power, and cabling of all disk shelves in the loop after the disk shelf reporting the autotermination.

#### Syslog Message

Disk shelf previously connected to the OUT port of ESH %s on channels %s disk shelf ID %d is now inaccessible. (Check cables, speed selection.)

#### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED

if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

## esh.term events

### esh.term.error

#### Severity

ERROR

#### Description

This message occurs when there is an error related to loop termination.

#### Corrective Action

Set the termination switch to the OFF position on the ESH module that is reporting the error.

#### Syslog Message

FORCED TERMINATION ERROR: TERM switch set improperly on ESH %s on channels %s disk shelf ID %d. The disk shelf connected to the OUT port of this ESH module is inaccessible.

#### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

## esh.unable events

### esh.unable.add.shelf

#### Severity

ERROR

#### Description

This message occurs when a disk shelf cannot be added due to a memory allocation failure during disk shelf discovery processing.

#### Corrective Action

If the error occurs upon addition of a disk shelf to a running system, disconnect the shelf and reattach it. If the error occurs during boot, reboot the controller.

#### Syslog Message

Could not perform disk shelf add because of memory allocation failure (%s).

#### Parameters

**name** (STRING): Where in the disk shelf add path the error occurred.

## esh.unexpected events

## esh.unexpected.status

### Severity

ALERT

### Description

This message occurs when a disk shelf reports an unknown status.

### Corrective Action

Ensure that the shelf module firmware and ONTAP® software are compatible versions. Find firmware versions at the NetApp technical support web site.

### Syslog Message

Unexpected ESH status of: %d.

### Parameters

**status** (INT): Unexpected/unknown status code that was received.

## esh.xfer.events

### esh.xfer.element.fault

### Severity

ERROR

### Description

This message occurs when an element is unreachable during an I/O request. It might be due to a transient condition in loop connectivity.

### Corrective Action

Check the cabling to the disk shelf. Verify that the SCSI Enclosure Services (SES) disk drives (0-1) are properly inserted. Check the status LED on the disk shelf, making sure that the fault LEDs are not lit. Clear any fault condition, if possible. Refer to the quick reference card beneath the disk shelf for information about the meanings of the LEDs.

### Syslog Message

I/O failure on ESH module %s on channels %s disk shelf id %d, status is %s.

### Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

**state** (STRING): State of ESH module.

### esh.xfer.not.sent

### Severity

ERROR

## Description

This message occurs when an I/O transfer cannot be sent. It might be due to a transient condition in loop connectivity.

## Corrective Action

Check the cabling to the disk shelf. Verify that the SCSI Enclosure Services (SES) disk drives (0-1) are properly inserted. Check the status LED on the disk shelf, making sure that the fault LEDs are not lit. Clear any fault condition, if possible. Refer to the quick reference card beneath the disk shelf for information about the meanings of the LEDs.

## Syslog Message

ESH %s I/O failure NOT\_SENT channels %s disk shelf ID %d.

## Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

## esh.xfer.pctl.error

### Severity

ERROR

### Description

This message occurs when an error is encountered while sending a disk 'power control' command to SCSI Enclosure Services (SES).

### Corrective Action

Check the cabling to the disk shelf. Verify that the SES disk drives in slots 0 and 1 are properly seated. Check the status LEDs on the disk shelf and make sure that the fault LEDs are not lit. Clear any fault condition, if possible. See the quick reference cards beneath the fault disk shelf for information about interpreting the LEDs.

### Syslog Message

Error sending disk 'power control' command to SCSI Enclosure Services (SES) in disk shelf %d on channel %s.

### Parameters

**shelfID** (INT): Disk shelf ID of the disk shelf that reported an error.

**channelName** (STRING): Name of the Fibre Channel adapter port.

## esh.xfer.unknown.error

### Severity

ERROR

### Description

This message occurs when there is an unknown error during an I/O request.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

ESH %s I/O failure UNKNOWN: %d channels %s disk shelf ID %d.

## Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**status** (INT): Error status code.

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.

## esh.xfer.xport.error

### Severity

ERROR

### Description

This message occurs when there is a transport error during an I/O request. It might be due to a transient condition in loop connectivity.

### Corrective Action

Check the cabling to the disk shelf. Verify that the SCSI Enclosure Services (SES) disk drives (0-1) are properly inserted. Check the status LED on the disk shelf, making sure that the fault LEDs are not lit. Clear any fault condition, if possible. Refer to the quick reference card beneath the disk shelf for information about the meanings of the LEDs.

## Syslog Message

ESH %s I/O failure TRANSPORT\_ERROR channels %s disk shelf ID %d.

## Parameters

**element** (STRING): ESH module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID (1-7) of the disk shelf reporting the error.



# exports events

## exports.anon events

### exports.anon.noCredForId

#### Severity

ERROR

#### Description

This message occurs when we cannot get UNIX credentials for a user id when evaluating the "-anon" portion of an export-policy rule.

#### Corrective Action

If "-anon" is a UID, ensure that a valid UID exists on the configured name servers. If "-anon" is a username, ensure that a mapping exists between the username and the UID on the configured name servers and the resulting UID is also present.

#### Syslog Message

Cannot retrieve credentials for UID "%u" present in the "-anon" entry for the export-policy rule for Vserver "%d".

#### Parameters

**anonId** (INT): UNIX UID for which we cannot acquire credentials.  
**vserver** (INT): Identifier of the Vserver owning the export-policy.

## exports.dns events

### exports.dns.config

#### Severity

ERROR

#### Description

This message occurs when a name services lookup request finds that DNS is configured as an ns-switch source for hosts, but DNS is not configured for the Vserver. Name services lookups using DNS do not function.

#### Corrective Action

Check the ns-switch sources for hosts database configured for the Vserver using "vserver services name-service ns-switch show -database hosts" and the DNS configuration for the Vserver using "vserver services name-service dns show". Either remove DNS as an ns-switch source, or configure DNS using the "vserver services name-service dns create" command.

#### Syslog Message

DNS is configured as a ns-switch source for hosts, but the DNS configuration is not found for Vserver "%s".

#### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

# exports.dom events

## exports.dom.notFound

### Severity

ERROR

### Description

This message occurs when the reverse lookup record mapping an IP address to its hostname cannot be found in the configured name servers. This record must be present to evaluate the export-policy rule that has a domain name in its clientmatch entry.

### Corrective Action

The reverse mapping record might have been incorrectly entered, inadvertently removed from the name service, or never entered at all. Check the configuration with the appropriate tools for each name service.

### Syslog Message

IP address "%s" does not have a reverse mapping for its corresponding hostname in the configured name servers when evaluating the export-policy rule at index "%u" in policy-id "%llu" for Vserver "%s".

### Parameters

**ipaddr** (STRING): IP address of the client that does not have the reverse lookup record.

**ruleindex** (INT): Index of the export-policy rule for which the problem is occurring.

**policyid** (LONGINT): Identifier for the export-policy owning the export-policy rule.

**vserver** (STRING): Name of the data Vserver for which the problem is occurring.

## exports.dom.transient

### Severity

ERROR

### Description

This message occurs when reverse lookup of the IP address of a client using the configured name servers in a data Vserver is unsuccessful for an export-policy rule that uses the domain name in clientmatch.

### Corrective Action

The name service might not be functioning correctly. Check the configuration with the appropriate tools for each name service.

### Syslog Message

IP address "%s" could not be resolved to a hostname (error %d) when evaluating export-policy rule at index "%u" in policy-id "%llu" in Vserver "%s".

### Parameters

**ipaddr** (STRING): IP address of the client encountering the reverse lookup error.

**retcode** (INT): Error code returned by the name service reverse lookup.

**ruleindex** (INT): Index of the export-policy rule for which the problem is occurring.

**policyid** (LONGINT): Identifier for the export-policy owning the export-policy rule.

**vserver** (STRING): Name of the data Vserver for which the problem is occurring.

# exports.hostname events

## exports.hostname.notFound

### Severity

INFORMATIONAL

### Description

This message occurs when the forward lookup record mapping a host name to its IP address is not found in the configured name servers in a Vserver. This record must be present to evaluate the export-policy rule that has a host name in its clientmatch entry.

### Corrective Action

Ensure that the host has an interface configured for this address family, that the host name is correctly entered, that the host was not inadvertently removed from the name service, or that the host name was never entered at all. If the host should have an address for this address family, then check the configuration by using the appropriate tools for each name service.

### Syslog Message

Hostname "%s" does not have a forward mapping for its corresponding %s address in the configured name servers when evaluating the export-policy rule at index "%u" in policy-id "%llu" for Vserver "%s".

### Parameters

**hostname** (STRING): Name of the host that is not found.

**addrfamily** (STRING): Name of the address family for which the problem occurs.

**ruleindex** (INT): Index of the export-policy rule for which the problem occurs.

**policyid** (LONGINT): Identifier for the export-policy that owns the export-policy rule.

**vserver** (STRING): Name of the Vserver for which the problem occurs.

## exports.hostname.transient

### Severity

ERROR

### Description

This message occurs when a host named in a clientmatch entry of an export-policy rule is not resolved to an IP address using the configured name servers in a data Vserver.

### Corrective Action

Ensure that the name service is functioning correctly by checking the configuration by using the appropriate tools for each name service.

### Syslog Message

Hostname "%s" could not be resolved to an %s address (error %d) when evaluating export-policy rule at index "%u" in policy-id "%llu" on Vserver "%s".

### Parameters

**hostname** (STRING): Name of the host that is not resolved.

**addrfamily** (STRING): Name of the address family for which the problem occurs.

**retcode** (INT): Error code returned by the name service lookup.

**ruleindex** (INT): Index of the export-policy rule for which the problem occurs.

**policyid** (LONGINT): Identifier for the export-policy that owns the export-policy rule.  
**vserver** (STRING): Name of the data Vserver for which the problem occurs.

## exports.netgroup events

### exports.netgroup.dnsNoPtrRec

#### Severity

ERROR

#### Description

This message occurs when the reverse lookup record mapping an IP address to its hostname cannot be found in the configured name servers. This record must be present to evaluate the membership of a hostname in the netgroup named in the clientmatch of an export-policy rule.

#### Corrective Action

The reverse mapping record might have been incorrectly entered, inadvertently removed from the name service, or never entered at all. Check the configuration with the appropriate tools for each name service.

#### Syslog Message

IP address "%s" does not have a reverse mapping for its corresponding hostname in the configured name servers when evaluating the export-policy rules for netgroup "%s" on Vserver "%s".

#### Parameters

**ipaddr** (STRING): IP address of the client that does not have the reverse lookup record.  
**netgroup** (STRING): Name of the netgroup that is being evaluated.  
**vserver** (STRING): Name of the Vserver for which the problem is occurring.

### exports.netgroup.notFound

#### Severity

ERROR

#### Description

This message occurs when a netgroup named in a clientmatch of an export-policy rule is not found in the configured name servers.

#### Corrective Action

The netgroup name might have been incorrectly entered, inadvertently removed from the name service, or never entered at all. Check the configuration with the appropriate tools for each name service.

#### Syslog Message

Netgroup "%s" cannot be found when evaluating export-policy rules on Vserver "%s".

#### Parameters

**netgroup** (STRING): Name of the netgroup that is not being found.  
**vserver** (STRING): Name of the Vserver for which the problem is occurring.

### exports.netgroup.partial

## Severity

ERROR

## Description

This message occurs when partial netgroup results are returned by the name services because of errors in the netgroup mapping.

## Corrective Action

If the netgroup mapping is greater than 1024 characters for netgroups on Network Information Service (NIS) or 4096 for netgroups in a local file, consider creating a sub netgroup to divide it. If the nesting of the netgroup is too deep, consider reducing the depth by adding the entries directly to the netgroup instead of adding to sub netgroups. If the netgroup has a sub netgroup that is not available as a top level netgroup in the mapping, add an empty entry for the sub netgroup to resolve the error. Ensure the line length and nesting limits are not exceeded.

## Syslog Message

Netgroup %s expansion failed with %s (error: %d) when evaluating export-policy rules on Vserver "%s".

## Parameters

**netgroup** (STRING): Name of the netgroup that could not be expanded.  
**reason** (STRING): Specific reason for the netgroup expansion failure.  
**error** (INT): Error code returned during netgroup expansion.  
**vserver** (STRING): Name of the Vserver in which the problem is occurring.

# exports.ngbh events

## exports.ngbh.allFailed

### Severity

ERROR

### Description

This message occurs when a netgroup by host request fails because all ns-switch sources for the netgroup database have returned connection errors and files are unusable as a source.

### Corrective Action

There might be temporary connectivity issues with the Vserver's configured ns-switch sources for the netgroup database. Check connectivity to the name servers configured for netgroups.

### Syslog Message

Connectivity issues encountered in all ns-switch sources when matching Client IP "%s" in Netgroup "%s" for Vserver "%s".

### Parameters

**client** (STRING): IP address of the client for which the lookup failed.  
**netgroup** (STRING): Name of the netgroup being evaluated.  
**vserver** (STRING): Name of the Vserver where the problem is occurring.

# exports.nsd events

## exports.nsd.anonNameTold

### Severity

ERROR

### Description

This message occurs when a user string cannot be mapped to a user id when evaluating the "-anon" portion of an export-policy rule.

### Corrective Action

If the "-anon" value is a name, ensure that a mapping from that name to a UID exists in the configured name services. For both a name and a UID entry in "-anon", ensure that the resulting UID exists in the configured name services.

### Syslog Message

Cannot retrieve UID for UNIX username "%s" present in the "-anon" entry for the export-policy rule for Vserver "%d".

### Parameters

**anon** (STRING): Anonymous string that cannot be mapped to a user id.  
**vserver** (INT): Name of the Vserver owning the export-policy.

# exports.policy events

## exports.policy.empty

### Severity

NOTICE

### Description

This message occurs when an export policy with no rules is assigned to a volume or a qtree. Such an empty policy makes the volume or qtree inaccessible to NFS and CIFS protocol clients.

### Corrective Action

If access is needed for the volume or qtree, then create an export policy rule using the "export-policy rule create" command.

### Syslog Message

An export policy "%s" with no rules was assigned to a %s "%s" on the Vserver %s.

### Parameters

**policyName** (STRING): Name of the export policy with no rules.  
**type** (STRING): Type of object (volume or qtree).  
**typeName** (STRING): Name of the volume or qtree.  
**vserver** (STRING): Name of the Vserver with this issue.

## **exports.policy.last.rule**

### **Severity**

ERROR

### **Description**

This message occurs when the last rule from the export policy assigned to a volume or qtree is deleted. An empty policy will make the volume or qtree inaccessible to clients using NFS or CIFS protocol.

### **Corrective Action**

If access is needed for the volume or qtree, then create an export policy rule using the "export-policy rule create" command.

### **Syslog Message**

The last rule from export policy "%s" was deleted on Vserver %s.

### **Parameters**

**policyName** (STRING): Name of the export policy from which the rule is being deleted.

**vserver** (STRING): Name of the Vserver where the issue occurs.

# extcache events

## extcache.aggregate events

### extCache.aggregate.invalidate

#### Severity

INFORMATIONAL

#### Description

This message provides diagnostic information about background work taking place in the WAFL® external cache subsystem.

#### Corrective Action

(None).

#### Syslog Message

WAFL external cache: %s

#### Parameters

**detail** (STRING): Specific diagnostic information related to the aggregate invalidation event.

## extcache.configunsupported events

### extCache.configUnsupported

#### Severity

ERROR

#### Description

This message occurs when the controller contains an external cache hardware configuration that cannot be supported on the controller.

#### Corrective Action

Consult the release documentation, check the Interoperability Matrix Tool (IMT) for a supported configuration, or contact Contact NetApp technical support. to understand and correct the hardware configuration.

#### Syslog Message

WAFL external cache hardware configuration is not supported: %s.

#### Parameters

**description** (STRING): String providing information about why the configuration is unsupported.

## extcache.control events

### extCache.control



**Severity**

NOTICE

**Description**

This message occurs when the external cache software cannot signal its administrative thread.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache: %s

**Parameters**

**detail** (STRING): Specific details of the event.

## extcache.disabled events

**extCache.disabled****Severity**

NOTICE

**Description**

This message occurs when the external cache is administratively disabled.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache administratively disabled

**Parameters**

(None).

## extcache.enabled events

**extCache.enabled****Severity**

NOTICE

**Description**

This message occurs when the external cache is administratively enabled.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache administratively enabled

## Parameters

(None).

# extcache.io events

## extCache.io.BlockChecksumError

### Severity

ERROR

### Description

This message occurs when the external cache detects a block checksum verification error while performing a read operation. Checksum errors are infrequent. Some combinations of external cache hardware and software handle these errors automatically without the need for administrator action. If further action is required, other messages will be generated. The operation will be retried from persistent storage (RAID).

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**block\_num** (INT): Block number that failed checksum verification.

## extCache.io.extInfo

### Severity

NOTICE

### Description

This message accompanies read or write errors and prints out any relevant information about the failure.

### Corrective Action

(None).

### Syslog Message

%s: fbn: %lld, pvbn: %lld, vvbn: %lld, inum: %d, ecbn: %lld, vol: %s, aggr: %s.

### Parameters

**description** (STRING): String providing information about the message.

**fileBlockNum** (LONGINT): fbn of the block that errored.

**physBlockNum** (LONGINT): pvbn of the block that errored.

**volBlockNum** (LONGINT): vvbn of the block that errored.

**fileId** (INT): Inode number of the file that owns the block.

**ECblockNum** (INT): The ecbn of the block.

**volName** (STRING): Name of the volume.

**aggrName** (STRING): Name of the aggregate.

## extCache.io.readError

### Severity

NOTICE

### Description

This message occurs when the external cache detects an error on a read. Read errors are infrequent and are usually handled automatically by the external cache without the need for administrator action. The operation will be retried from persistent storage (RAID).

### Corrective Action

(None).

### Syslog Message

WAFL external cache I/O read error: %s, code %d.

### Parameters

**description** (STRING): String providing information about why the operation failed.

**detail** (INT): Detailed error code about why the operation failed.

## extCache.io.writeError

### Severity

NOTICE

### Description

This message occurs when the external cache detects an error on a write. Write errors are infrequent and are usually handled automatically by the external cache without the need for administrator action.

### Corrective Action

(None).

### Syslog Message

WAFL external cache I/O write error: %s, code %d.

### Parameters

**description** (STRING): String providing information about why the operation failed.

**detail** (INT): Detailed error code about why the operation failed.

## extcache.offline events

### extCache.offline

### Severity

ERROR

### Description

This message occurs when the external cache is automatically taken offline. This can happen after an I/O error on the external cache and may result in degraded performance until the problem is corrected.

**Corrective Action**

(Call support).

**Syslog Message**

External cache has been taken offline on error (rc=%d). Check the EMS log for earlier errors.

**Parameters**

**code** (INT): A unique error code for this event.

## extcache.qos events

**extCache.QoS****Severity**

INFORMATIONAL

**Description**

This message occurs when External Cache Quality of Service (QoS) monitoring is enabled/disabled.

**Corrective Action**

(None).

**Syslog Message**

External Cache QoS is: %s.

**Parameters**

**description** (STRING): State of External Cache QoS

## extcache.restart events

**extCache.restart****Severity**

NOTICE

**Description**

This message occurs when the external cache encounters an internal consistency issue and triggers an automatic reset.

**Corrective Action**

(None).

**Syslog Message**

WAF external cache reset: %s

**Parameters**

**detail** (STRING): String providing additional information about what triggered the reset.

# extcache.rw events

## extCache.rw.alloc

### Severity

NOTICE

### Description

This message occurs when an external cache cannot obtain sufficient system resources. This causes the external cache component to be unable to rewarm after a reboot, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache could not allocate resources: %s.

### Parameters

**description** (STRING): String identifying the specific resources that were not obtained.

## extCache.rw.canceled

### Severity

NOTICE

### Description

This message occurs when the external cache cannot complete cache rearming. This causes the external cache component to start with no data cached, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache reconstruct was canceled.

### Parameters

(None).

## extCache.rw.configChange

### Severity

INFORMATIONAL

### Description

This message occurs when the external cache warming logic detects a change in cache hardware configuration. This causes the external cache component to start with no data cached for any aggregates, which can cause a temporary performance impact. The problem should resolve itself at the next reboot.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache warming aborted due to hardware configuration change.

**Parameters**

(None).

**extCache.rw.emptyCache****Severity**

INFORMATIONAL

**Description**

This message occurs when the external cache warming process terminates due to a lack of active cache entries.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache warming canceled: cache is not in use.

**Parameters**

(None).

**extCache.rw.entry****Severity**

NOTICE

**Description**

This message occurs when the external cache cannot rebuild cache entries from stored data. This event is caused by an internal error and results in a small portion of the cache remaining empty during rewarming.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache experienced an internal error during reconstruct.

**Parameters**

(None).

**extCache.rw.headerInvalid****Severity**

NOTICE

## Description

This message occurs when the external cache warming process encounters an invalid file header. This causes the external cache component to start with no data cached for one or all aggregates, which might cause a temporary performance impact.

## Corrective Action

(None).

## Syslog Message

WAFL external cache warming encountered an invalid file header: %s.

## Parameters

**section** (STRING): String identifying the header section in question.

## extCache.rw.io.readLimit

### Severity

NOTICE

### Description

This message occurs when the external cache exceeds the allowable threshold for I/O errors on reads. This causes the external cache component to start with no data cached for any aggregates and might result in temporarily degraded performance.

### Corrective Action

(None).

### Syslog Message

WAFL external cache warming encountered too many read errors: %s (%d).

### Parameters

**description** (STRING): String providing information about the operation.

**number** (INT): Number of unsuccessful reads detected.

## extCache.rw.io.write

### Severity

INFORMATIONAL

### Description

This message occurs when the external cache detects an I/O error on a write. This causes the external cache component to be unable to reconstruct portions of the cache and might result in slightly degraded performance.

### Corrective Action

(None).

### Syslog Message

WAFL external cache I/O write: %s, code %d.

## Parameters

**description** (STRING): String providing information about the operation.

**detail** (INT): Detailed error code about the operation.

## extCache.rw.io.writeLimit

### Severity

NOTICE

### Description

This message occurs when the external cache exceeds the allowable threshold for I/O errors on writes.

This causes the external cache component to start with no cached data for any aggregates and might result in temporarily degraded performance.

### Corrective Action

(None).

### Syslog Message

WAFL external cache warning encountered too many write errors: %s (%d).

## Parameters

**description** (STRING): String providing information about the operation.

**number** (INT): Number of unsuccessful writes detected.

## extCache.rw.log.bufLoad

### Severity

NOTICE

### Description

This message occurs when the external cache invalidation log cannot load a WAFL® buffer. This causes the external cache component to be unable to rewarm the aggregate in question after a reboot, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache invalidation log buf could not be loaded: aggregate %s.

## Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

## extCache.rw.log.disabled

### Severity

NOTICE

### Description

This message occurs when external cache invalidation logging is disabled for an aggregate. This causes the external cache component to be unable to rewarm that aggregate after a reboot, which might cause a



temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache invalidation log disabled: aggregate %s: %s.

### Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**reason** (STRING): String indicating the cause of the problem.

## extCache.rw.log.freeSpace

### Severity

NOTICE

### Description

This message occurs when an external cache log file cannot be written due to lack of free space. This causes the external cache component to be unable to rewarm a specific aggregate or the entire cache after a reboot, which might cause a temporary performance impact.

### Corrective Action

Make sufficient space available on the aggregate.

### Syslog Message

WAFL external cache log out of space on aggregate %s (needs %d bytes).

### Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the failure.

**bytes** (INT): Number of free bytes needed to successfully open the log file.

## extCache.rw.log.maxSize

### Severity

NOTICE

### Description

This message occurs when an external cache log file reaches its maximum size. This causes the external cache component to be unable to rewarm a specific aggregate or the entire cache after a reboot, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache log reached maximum size: aggregate %s, log %s.

### Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**logtype** (STRING): String identifying the log file that reached the maximum size.

## extCache.rw.log.open

### Severity

NOTICE

### Description

This message occurs when the external cache cannot open a log file. This causes the external cache component to be unable to rewarm some or all of the cached aggregates after a reboot, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache log could not be opened: aggregate %s, log %s.

### Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**logtype** (STRING): String identifying the type of log file that could not be opened.

## extCache.rw.log.range

### Severity

NOTICE

### Description

This message occurs when the external cache invalidation log has a block address that is not within an aggregate's address range. This causes the external cache component to be unable to rewarm the aggregate in question after a reboot, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache invalidation log block out of range: aggregate %s block %lld.

### Parameters

**aggregate** (STRING): String identifying the aggregate that experienced the failure.

**block** (LONGINT): Block encountered in the invalidation log.

## extCache.rw.log.read

### Severity

NOTICE

### Description

This message occurs when the external cache cannot read from a log file. This causes the external cache component to be unable to rewarm some or all of the cached aggregates during a reboot, which might cause a temporary performance impact.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache log could not be read: aggregate %s, log %s.

**Parameters**

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**logtype** (STRING): String identifying the log file that could not be read.

**extCache.rw.log.seek****Severity**

NOTICE

**Description**

This message occurs when the external cache cannot manipulate a log file. This causes the external cache component to be unable to rewarm after a reboot, which might cause a temporary performance impact.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache log could not be read: aggregate %s, log %s.

**Parameters**

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**logtype** (STRING): String identifying the log file that could not be read.

**extCache.rw.log.write****Severity**

NOTICE

**Description**

This message occurs when the external cache cannot write to a log file. This causes the external cache component to be unable to rewarm some or all of the cached aggregates after a reboot, which might cause a temporary performance impact.

**Corrective Action**

(None).

**Syslog Message**

WAFL external cache log could not be written: aggregate %s, log %s.

**Parameters**

**aggregate** (STRING): String identifying the aggregate that experienced the problem.

**logtype** (STRING): String identifying the log file that could not be written.

## **extCache.rw.replay.canceled**

### **Severity**

NOTICE

### **Description**

This message occurs when the external cache cannot successfully process the invalidation logs for an individual aggregate. This causes the external cache component to start with no data cached for the aggregate, which might cause a temporary performance impact.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL external cache replay canceled for aggregate %s: %s.

### **Parameters**

**aggregate** (STRING): String identifying the aggregate in question.

**description** (STRING): String providing information about the operation.

## **extCache.rw.replay.complete**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the external cache finishes processing the invalidation logs for an individual aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL external cache warming replay complete for aggregate %s with invalidation log %d%% full.

### **Parameters**

**aggregate** (STRING): String identifying the aggregate in question.

**percent** (INT): Percentage of ilog that is full.

## **extCache.rw.replay.start**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the external cache begins processing the invalidation logs for an individual aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL external cache warming replay has begun for aggregate %s.

### **Parameters**

**aggregate** (STRING): String identifying the aggregate in question.

## **extCache.rw.replay.timeout**

### **Severity**

NOTICE

### **Description**

This message occurs when the external cache timed out waiting for an aggregate to come online and purged old cache entries for it. This causes the external cache component to start with no data cached for the aggregate, which might cause a temporary performance impact.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL external cache replay timed out for aggregate %s.

### **Parameters**

**aggregate** (STRING): String identifying the aggregate in question.

## **extCache.rw.snap.canceled**

### **Severity**

NOTICE

### **Description**

This message occurs when the external cache cannot save a warming snapshot copy. This causes the external cache component to be unable to rewarm after a reboot, which might cause a temporary performance impact.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL external cache snapshot copy canceled: %s.

### **Parameters**

**description** (STRING): String providing information about why the operation was canceled.

## **extCache.rw.terminated**

### **Severity**

NOTICE

## Description

This message occurs when the external cache warming process is terminated. This causes the external cache component to start with no data cached for any aggregates, which might cause a temporary performance impact.

## Corrective Action

(None).

## Syslog Message

WAFL external cache warming process terminated.

## Parameters

(None).

## extCache.rw.timeout

### Severity

NOTICE

### Description

This message occurs when the external cache warming process terminates after exceeding its allowed time. This causes the external cache component to start with a subset of cached data, which might cause a temporary performance impact.

### Corrective Action

(None).

### Syslog Message

WAFL external cache warming process exceeded its maximum allowable time and was aborted.

### Parameters

(None).

## extcache.unallowed events

### extCache.unallowed.option

#### Severity

NOTICE

#### Description

This message occurs when an external cache bootarg or setflag is set without the required privilege. The default value is used instead.

#### Corrective Action

(None).

#### Syslog Message

WAFL external cache %s %s is set without the required privilege; using the default value.

## Parameters

**type** (STRING): Whether the option set is a bootarg or a setflag.

**option** (STRING): Option name.

# fabric events

## fabric.mgmt events

### fabric.mgmt.interface

**Severity**

NOTICE

**Description**

This message occurs when the system fails to establish a management interface on the fabric.

**Corrective Action**

(None).

**Syslog Message**

Cannot establish a management (Status/Control/Maintenance) interface on fabric %s.

**Parameters**

**fabName** (STRING): Fabric name.

## fabric.sdm events

### fabric.sdm.mem.alloc

**Severity**

NOTICE

**Description**

This message occurs when the system fails to allocate a message for storage device monitor (SDM) notification.

**Corrective Action**

(None).

**Syslog Message**

Failed to allocate a message for SDM notification to fabric\_man.

**Parameters**

(None).



# fabriclink events

## fabriclink.retry events

### fabriclink.retry.delay

#### Severity

NOTICE

#### Description

This message occurs when object replication encounters an error when attempting to contact a remote object store. The attempt will be retried after a short delay.

#### Corrective Action

(None).

#### Syslog Message

Link L%u (%s) encountered OSC error %s; will retry in %u seconds.

#### Parameters

**linkid** (INT): Link ID that encountered an error.

**objstore** (STRING): Object store UUID that encountered an error.

**error** (STRING): Error that was encountered.

**delay** (INT): Delay (seconds) before the attempt will be retried.

# fabricpool events

## fabricpool.cloud events

### fabricpool.cloud.subscription.added

**Severity**

NOTICE

**Description**

This message occurs when a cloud subscription license for FabricPools has been added.

**Corrective Action**

(None).

**Syslog Message**

A cloud subscription license for FabricPools has been added for license type %s. Cluster id: %s.

**Parameters**

**license\_type** (STRING): The type of license that the subscription is added for.

**cluster\_uuid** (STRING): UUID of the cluster.

### fabricpool.cloud.subscription.removed

**Severity**

NOTICE

**Description**

This message occurs when a cloud subscription license for FabricPools has been removed.

**Corrective Action**

(None).

**Syslog Message**

A cloud subscription license for FabricPools has been removed for license type %s. Cluster id: %s.

**Parameters**

**license\_type** (STRING): The type of license that the subscription was removed for.

**cluster\_uuid** (STRING): UUID of the cluster.

## fabricpool.full events

### fabricpool.full

**Severity**

ERROR

## Description

This message occurs when the total, cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the license limit.

## Corrective Action

To find the percentage of the licensed capacity used by each FabricPool, use the "storage aggregate object-store show-space" command. To create more space, delete Snapshot(tm) copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command. To increase the licensed capacity, install a new license on the cluster.

## Syslog Message

Total, cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the licensed limit. Cluster ID: %s. Current usage: %s, licensed capacity: %s.

## Parameters

**cluster\_uuid** (STRING): UUID of the cluster.

**used\_capacity** (STRING): Total usage of capacity-licensed object stores.

**licensed\_capacity** (STRING): FabricPool licensed capacity.

# fabricpool.lost events

## fabricpool.lost.contact.scu

### Severity

ALERT

### Description

This message occurs when ONTAP has lost contact with the Service Control Unit used with cloud subscription licenses.

### Corrective Action

Ensure that the Service Control Unit is running and Ontap should have connectivity to it.

### Syslog Message

ONTAP has lost contact with the service control unit for the license of type %s. Cluster ID: %s.

### Parameters

**license\_type** (STRING): The type of license which has lost contact with the Service Control Unit.

**cluster\_uuid** (STRING): UUID of the cluster.

# fabricpool.nearly events

## fabricpool.nearly.full

### Severity

ERROR

### Description

This message occurs when the total, cluster-wide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the license limit.

## Corrective Action

To find the percentage of the licensed capacity used by each FabricPool, use the "storage aggregate object-store show-space" command. To create more space, delete Snapshot® copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command. To increase the licensed capacity, install a new license on the cluster.

## Syslog Message

Total, cluster-wide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the licensed limit. Cluster id: %s. Current usage: %s, licensed capacity: %s.

## Parameters

**cluster\_uuid** (STRING): UUID of the cluster.

**used\_capacity** (STRING): Total usage of capacity-licensed object stores.

**licensed\_capacity** (STRING): FabricPool licensed capacity.

# fabricpool.scu events

## fabricpool.scu.contacted

### Severity

NOTICE

### Description

This message occurs when contact has resumed between ONTAP and the Service Control Unit that is used with cloud subscription licenses.

### Corrective Action

(None).

### Syslog Message

Contact has resumed between ONTAP and the service control unit for the license of type %s. Cluster ID: %s.

### Parameters

**license\_type** (STRING): The type of license which has resumed contact with the Service Control Unit.

**cluster\_uuid** (STRING): UUID of the cluster.

# fabricpool.tiering events

## fabricpool.tiering.stopped.no.scu

### Severity

ALERT

### Description

This message occurs when ONTAP has stopped tiering data to the cloud tier because it has lost contact with the Service Control Unit used for one or more cloud subscription licenses.

### Corrective Action

Ensure that the Service Control Unit is running and ONTAP should have connectivity to it.

## **Syslog Message**

Ontap has stopped tiering data to the cloud tier because it lost contact with the Service Control Unit for the license of type %s. Cluster ID: %s.

## **Parameters**

**license\_type** (STRING): The type of license which has lost contact with the Service Control Unit.

**cluster\_uuid** (STRING): UUID of the cluster.

# **fal\_nvme events**

## **fal\_nvme.init events**

### **fal\_nvme.init.failed**

**Severity**

ERROR

**Description**

This message occurs when the Flash Adaptation Layer (FAL) NVMe subsystem fails to initialize.

**Corrective Action**

Reboot the node. If the problem persists, contact technical support for assistance.

**Syslog Message**

FAL NVMe subsystem initialization failed. Reason: %s.

**Parameters**

**reason** (STRING): Reason for failure.

## **fal\_nvme.online events**

### **fal\_nvme.online.succ**

**Severity**

INFORMATIONAL

**Description**

This message occurs when the Flash Adaptation Layer (FAL) NVMe subsystem successfully brings online all resources in a caching module.

**Corrective Action**

(None).

**Syslog Message**

Successfully brought online %lu GB of the caching module.

**Parameters**

**capacity** (LONGINT): Total capacity.

## **fal\_nvme.partition events**

### **fal\_nvme.partition.status**

**Severity**

NOTICE

**Description**

This message occurs when a partition status in the Flash Adaptation Layer (FAL) changes.

**Corrective Action**

(None).

**Syslog Message**

Partition %lx-%lu with capacity %lu GiB status: %s.

**Parameters**

**id** (LONGINTEX): Unique identifier for the partition.

**ns** (LONGINTEX): Index of the namespace.

**capacity** (LONGINT): Partition capacity.

**state** (STRING): Partition state.

# fc events

## fc.bridge events

### fc.bridge.fc.port.alert

#### Severity

ALERT

#### Description

This message occurs when an FC port goes offline.

#### Corrective Action

Check cables, SFPs, and connected devices. If the port is intentionally unused, use the Fibre Channel bridge CLI to disable it.

#### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s FC Port FC\_PORT%d offline. %s

#### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePort** (INT): Port reported by the FC-to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

### fc.bridge.fc.port.notice

#### Severity

NOTICE

#### Description

This message occurs when an FC port on the identified bridge comes online.

#### Corrective Action

(None).

#### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s FC Port FC\_PORT%d connected. %s

#### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePort** (INT): Port reported by the FC-to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.



## fc.bridge.power.alert

### Severity

ALERT

### Description

This message occurs when an FC-to-SAS bridge power supply fails.

### Corrective Action

Check input power and cable. If the problem persists, replace the power supply.

### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s Power Supply %s has failed. %s

### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePowerSupply** (STRING): Power supply reported by the FC-to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.power.notice

### Severity

NOTICE

### Description

This message occurs when an FC-to-SAS bridge power supply returns to normal operation.

### Corrective Action

(None).

### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s Power supply %s operating normally. %s

### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePowerSupply** (STRING): Port reported by the FC-to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.sas.phy.alert

### Severity

ALERT

### Description

This message occurs when a SAS port PHY goes offline.

## Corrective Action

Check Mini-SAS HD connector, cables, and connected equipment.

## Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s SAS Port %d Phy %d offline. %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePort** (INT): Port reported by the FC-to-SAS bridge.

**bridgePhy** (INT): SAS port Phy reported by the FC-to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.sas.phy.notice

### Severity

NOTICE

### Description

This message occurs when a SAS port PHY on a bridge comes online.

### Corrective Action

(None).

### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s SAS port %d Phy %d online. %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.

**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.

**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.

**bridgePort** (INT): Port reported by the FC-to-SAS bridge.

**bridgePhy** (INT): Phy reported by the FC—to-SAS bridge.

**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.sas.port.alert

### Severity

ALERT

### Description

This message occurs when an SAS port goes offline.

### Corrective Action

Check cables, SFPs, and connected devices. If the port is intentionally unused, use the Fibre Channel bridge CLI to disable it.

## Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s SAS Port SAS\_PORT%d offline. %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**bridgePort** (INT): Port reported by the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.sas.port.notice

### Severity

NOTICE

### Description

This message occurs when an SAS port on the identified bridge comes online.

### Corrective Action

(None).

## Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s SAS Port SAS\_PORT%d connected. %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**bridgePort** (INT): Port reported by the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.temp.alert

### Severity

ALERT

### Description

This message occurs when an FC-to-SAS bridge reaches an elevated temperature threshold. The bridge will continue to operate, but action should be taken to reduce the temperature.

### Corrective Action

The bridge is experiencing an elevated temperature, which indicates a loss of cooling airflow. Check cooling airflow, fan operation, and ambient temperature, and take action to reduce the FC-to-SAS bridge operating temperature. If this is not done and the temperature continues to rise, the FC-to-SAS bridge will shut down when the critical temperature threshold is reached.

## Syslog Message

FC-to-SAS %s bridge with name: %s S/N: %s WWN: %s is operating at elevated temperature (%d C). %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**bridgeTemp** (INT): Temperature reported by the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.temp.emergency

### Severity

EMERGENCY

### Description

This message occurs when an FC-to-SAS bridge reaches a critical temperature threshold. The bridge will shut down to protect the hardware.

### Corrective Action

Take immediate action to reduce the operating temperature of the FC-to-SAS bridge to avoid hardware damage. Check cooling airflow, fan operation, and ambient temperature. The FC-to-SAS bridge will shut down when this condition is reached, and will not operate until the temperature is reduced.

### Syslog Message

FC-to-SAS %s bridge with name: %s S/N: %s WWN:%s has reached critical temperature (%d C) and will soon shut down. %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**bridgeTemp** (INT): Temperature reported by the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.temp.notice

### Severity

NOTICE

### Description

This message occurs when an FC-to-SAS bridge returns to a normal operating temperature. No action needs to be taken, because the FC-to-SAS bridge temperature has returned to normal levels.

### Corrective Action

(None).

### Syslog Message

FC-to-SAS %s bridge %s S/N: %s WWN: %s temperature has returned to normal (%d C). %s

## Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.

**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**bridgeTemp** (INT): Temperature reported by the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.throughput.alert

### Severity

ALERT

### Description

This message occurs when an FC-to-SAS bridge experiences degraded throughput.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s is experiencing degraded throughput. %s

### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

## fc.bridge.throughput.notice

### Severity

NOTICE

### Description

This message occurs when an FC-to-SAS bridge recovers from degraded throughput.

### Corrective Action

(None).

### Syslog Message

Fibre Channel %s Bridge %s S/N: %s WWN: %s throughput is normal. %s

### Parameters

**bridgeType** (STRING): Type of the FC-to-SAS bridge.  
**bridgeName** (STRING): Name of the FC-to-SAS bridge.  
**bridgeSerialNum** (STRING): Serial number of the FC-to-SAS bridge.  
**bridgeWWN** (STRING): World-Wide Name (WWN) of the FC-to-SAS bridge.  
**addlInfo** (STRING): Additional information about the alert being generated, if available.

# fcadmin events

## fcadmin.bridge events

### fcadmin.bridge.prog.disk

#### Severity

INFORMATIONAL

#### Description

This event reports the WWN written to a device.

#### Corrective Action

(None).

#### Syslog Message

Programming Disk %s LUN=%d on Path=0x%0x with WWN  
%02x:%02x:%02x:%02x:%02x:%02x:%02x:%02x.

#### Parameters

**diskName** (STRING): Name of the disk.  
**LUNid** (INT): Logical Unit Number identifier.  
**pathAddress** (INTHEX): Device address.  
**WWNbyte0** (INTHEX): 0th byte of World Wide Name.  
**WWNbyte1** (INTHEX): 1st byte of World Wide Name.  
**WWNbyte2** (INTHEX): 2nd byte of World Wide Name.  
**WWNbyte3** (INTHEX): 3rd byte of World Wide Name.  
**WWNbyte4** (INTHEX): 4th byte of World Wide Name.  
**WWNbyte5** (INTHEX): 5th byte of World Wide Name.  
**WWNbyte6** (INTHEX): 6th byte of World Wide Name.  
**WWNbyte7** (INTHEX): 7th byte of World Wide Name.

### fcadmin.bridge.prog.done

#### Severity

INFORMATIONAL

#### Description

This message indicates that WWN programming has been successfully completed on a disk.

#### Corrective Action

(None).

#### Syslog Message

Disk %s LUN=%d on path=0x%0x successfully programmed. Reboot the drive to see the updated WWN.

#### Parameters

**diskName** (STRING): Name of the disk.  
**LUNid** (INT): Logical Unit Number identifier.  
**pathAddress** (INTHEX): Device address.

## fcadmin.bridge.prog.fault

### Severity

ERROR

### Description

This message indicates that the attempt to write the WWN to the disk has failed.

### Corrective Action

Retry the command after sometime, error might be transient.

### Syslog Message

Error programming disk %s LUN=%d on path=0x%0x. Error status=0x%0x.

### Parameters

**diskName** (STRING): Name of the disk.  
**LUNid** (INT): Logical Unit Number identifier.  
**pathAddress** (INTHEX): Device address.  
**commandStatus** (INTHEX): Returned SCSI error code.

## fcadmin.bridge.progWWN

### Severity

INFORMATIONAL

### Description

This message reports the statistics of the operation to write the World Wide Name (WWN) to the drives.

### Corrective Action

(None).

### Syslog Message

Total LUNs/drives reported = %d. Total LUNs/drives programmed = %d.

### Parameters

**reportedLUNs** (INT): Total number of LUN instances reported for entire session.  
**programmedLUNs** (INT): Total number of LUN instances programmed for entire session.

## fcadmin.bridge.retrieve.fail

### Severity

ERROR

### Description

This event is generated when the drive is not qualified for DDB (Disk Drive Behavior) update while programming the WWNs.

### Corrective Action

(None).

### **Syslog Message**

Error retrieving DDB data for disk=%s on path=0x%0x. Drives on the current path could not be programmed with WWN.

### **Parameters**

**diskName** (STRING): Name of the disk.  
**pathAddress** (INTHEX): Device address.

## **fcadmin.bridge.rprt.fault**

### **Severity**

ERROR

### **Description**

This message is generated when the Report LUNs response completed with some error.

### **Corrective Action**

Retry the command after sometime, error might be transient.

### **Syslog Message**

Error reported for Report LUN command on disk %s on path 0x%0x, error=0x%0x.

### **Parameters**

**diskName** (STRING): Name of the disk.  
**pathAddress** (INTHEX): Device address.  
**commandStatus** (INTHEX): Returned SCSI error code.

## **fcadmin.bridge.rprtcmd.fail**

### **Severity**

ERROR

### **Description**

This event is generated when insufficient SCSI command blocks are available for Report LUN command.

### **Corrective Action**

Retry the command after sometime. As command blocks may have become available.

### **Syslog Message**

Memory allocation failed for Report LUN command. No LUN instances are programmed.

### **Parameters**

(None).

## **fcadmin.bridge.rprtdata.fail**

### **Severity**

ERROR



**Description**

This event is generated when insufficient SCSI command blocks are available for Report LUN command.

**Corrective Action**

Retry the command after sometime. As command blocks may have become available.

**Syslog Message**

Memory allocation failed for Report LUN data command, bytes Requested=%d Disk %s on Path=0x%0x.

**Parameters**

**requestedMemory** (INT): Requested memory for Report LUN data command.

**diskName** (STRING): Name of the disk

**deviceAddress** (INTHEX): Device address.

**fcadmin.bridge.WWN.prog.stat****Severity**

INFORMATIONAL

**Description**

This event is generated when the programming of the WWNs into the reported LUNs is complete. An error in programming is indicated when the number of reported LUNs is different from the number programmed.

**Corrective Action**

(None).

**Syslog Message**

WWN update completed on disk %s path 0x%0x, total %d LUNS reported, total %d LUNS programmed.

**Parameters**

**diskName** (STRING): Name of the disk.

**pathAddress** (INTHEX): Device address.

**LUNsReported** (INT): Number of LUNs reported.

**LUNsProgrammed** (INT): Number of LUNs programmed.

# fci events

## fci.adapter events

### fci.adapter.bootarg.option

#### Severity

INFORMATIONAL

#### Description

This message occurs when the the Fibre Channel adapter driver is setting an option based on the setting of a bootarg variable.

#### Corrective Action

(None).

#### Syslog Message

Bootarg option "%s" configured for Fibre Channel adapter %s.

#### Parameters

**option** (STRING): Option being modified.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter (HBA).

### fci.adapter.dump.saved

#### Severity

NOTICE

#### Description

This message occurs when an unexpected event or invalid condition is detected by the Fibre Channel(FC) initiator driver and the contents of the FC initiator memory is saved. Dump files are stored in a directory on the system's root volume at /etc/log, with the following directory name format:  
slifc\_[adapterName][YYYYMMDD][HHMMSS]

#### Corrective Action

Send the dump file to technical support for analysis.

#### Syslog Message

A dump for adapter %s was stored in %s.

#### Parameters

**adapterName** (STRING): Name of the associated FC host bus adapter.

**filename** (STRING): Name of the file containing the dump.

### fci.adapter.error

#### Severity

NOTICE

## Description

This message occurs when the Fibre Channel adapter driver encounters an error with the adapter. The adapter will be reset to recover.

## Corrective Action

(None).

## Syslog Message

Fibre Channel adapter driver encountered error "%s" on adapter %s.

## Parameters

**error** (STRING): Type of error that occurred.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.fabric.login.failed

### Severity

ERROR

### Description

This message occurs when the Fibre Channel adapter cannot log in with the fabric management server.

### Corrective Action

The driver will attempt to recover from this error, but if this error keeps occurring, make sure that the switch is configured correctly and then try disconnecting and then reconnecting the adapter to the switch.

### Syslog Message

Fibre Channel adapter %s cannot log in with the fabric management server for reason 0x%x.

### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**status** (INHEX): Status returned from the adapter.

## fci.adapter.fabric.ns.login.failed

### Severity

ERROR

### Description

This message occurs when the Fibre Channel adapter cannot log in with the fabric name server.

### Corrective Action

The driver will attempt to recover from this error but if this error keeps occurring make sure that the switch is configured correctly, and then try disconnecting and then reconnecting the adapter to the switch.

### Syslog Message

Fibre Channel adapter %s was unable to log in with the fabric name server for reason 0x%x.

### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**status** (INHEX): Status returned from the adapter.

## **fci.adapter.fabric.request.failed**

### **Severity**

ERROR

### **Description**

This message occurs when the Fibre Channel adapter returns an unexpected status for a fabric query.

### **Corrective Action**

The driver will attempt to recover from this error but if this error keeps occurring make sure that the switch is configured correctly, and then try disconnecting and then reconnecting the adapter to the switch.

### **Syslog Message**

Fibre Channel adapter %s returned unexpected status 0x%x for a fabric query.

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**status** (INHEX): Status returned from the adapter.

## **fci.adapter.failed**

### **Severity**

ALERT

### **Description**

This message occurs when the Fibre Channel adapter driver cannot recover the adapter after resetting it multiple times. The adapter will be put offline.

### **Corrective Action**

If the adapter is in use check the cabling, and if it is connected to disk shelves check the seating of the LRC/ESH cards and disks. If the problem keeps occurring try replacing the adapter. If the issue is still not resolved, contact NetApp technical support for assistance in diagnosing a Fibre Channel issue.

### **Syslog Message**

Failed to recover Fibre Channel HBA %s after multiple resets. Taking the adapter offline.

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**reason** (STRING): Diagnostic information.

## **fci.adapter.firmware.fault**

### **Severity**

NOTICE

### **Description**

This message occurs when a firmware fault is detected on the Fibre Channel adapter and it is being reset to recover.

### **Corrective Action**

(None).

## Syslog Message

Detected firmware fault 0x%x on Fibre Channel adapter %s.

## Parameters

**faultCode** (INHEX): Firmware fault code.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.firmware.pause

### Severity

NOTICE

### Description

This message occurs when a firmware fault is detected on the Fibre Channel adapter and it is being reset to recover.

### Corrective Action

(None).

## Syslog Message

Detected firmware fault 0x%x on Fibre Channel adapter %s.

## Parameters

**faultCode** (INHEX): Firmware fault code.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.link.online

### Severity

INFORMATIONAL

### Description

This message occurs when a Fibre Channel adapter reports online link status, usually after the adapter is reset successfully.

### Corrective Action

(None).

## Syslog Message

Fibre Channel adapter %s link online.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.notification

### Severity

NOTICE

**Description**

This message occurs following a management request from the Fibre Channel(FC) driver to the host bus adapter.

**Corrective Action**

(None).

**Syslog Message**

Fibre Channel adapter driver notification: "%s" on adapter %s.

**Parameters**

**info** (STRING): Description of the outcome.

**adapterName** (STRING): Name of the associated FC host bus adapter.

**fci.adapter.nvram.invalid****Severity**

ALERT

**Description**

This message occurs when ONTAP® software cannot read the NVRAM content of a Fibre Channel adapter. This might cause ONTAP to load an incorrect driver, which could cause the adapter not to work as intended. In addition, the World Wide Node Name (WWNN) and the World Wide Port Name (WWPN) will be incorrect and temporary ones will be assigned. This might cause fabric zoning not to work correctly and might cause problems with anything that depends on consistent WWNNs and WWPNs.

**Corrective Action**

Reboot the appliance. If the problem still exists, replace the adapter.

**Syslog Message**

Cannot read NVRAM content of Fibre Channel adapter %s.

**Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**fci.adapter.offline****Severity**

INFORMATIONAL

**Description**

This message occurs when the system takes a Fibre Channel adapter offline.

**Corrective Action**

(None).

**Syslog Message**

Fibre Channel adapter %s is now offline.

**Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## **fci.adapter.offlining**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the system takes a Fibre Channel adapter offline after all outstanding I/O requests have finished.

### **Corrective Action**

(None).

### **Syslog Message**

Taking Fibre Channel adapter %s offline.

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## **fci.adapter.online**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the system brings a Fibre Channel adapter online.

### **Corrective Action**

(None).

### **Syslog Message**

Fibre Channel adapter %s is now online.

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## **fci.adapter.online.failed**

### **Severity**

ERROR

### **Description**

This message occurs when a Fibre Channel adapter fails to come online.

### **Corrective Action**

If the adapter is in use check the cabling, and if it connected to disk shelves check the seating of the LRC/ESH cards.

### **Syslog Message**

Fibre Channel adapter %s failed to come online.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.onlining

### Severity

INFORMATIONAL

### Description

This message occurs when the Fibre Channel adapter begins the process of going online.

### Corrective Action

(None).

### Syslog Message

Onlining Fibre Channel adapter %s.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.adapter.speed

### Severity

INFORMATIONAL

### Description

This message occurs when speed auto- negotiation for the Fibre Channel port is disabled because the administrator used either the 'ispfc-force-XG-only?' or the 'mptfc-force-XG-only?' boot argument in conjunction with the 'fc-force-XG-ports' boot argument to force a particular speed. For example, setting 'ispfc-force-2G-only?' to 'true' and setting 'fc-force-2G-ports' to '0a,0c' would force ports 0a and 0c to operate at a rate of 2 Gigabits per second.

### Corrective Action

(None).

### Syslog Message

The speed of Fibre Channel adapter port %s was forced to %u Gbps by link speed related boot arguments.

## Parameters

**port** (STRING): Name of the adapter port.

**speed** (INT): Speed the adapter port was assigned.

## fci.adapter.topology

### Severity

INFORMATIONAL

### Description

This message occurs when topology auto- negotiation for the Fibre Channel port is disabled because the administrator used either the 'fc-force-ptp-ports' or the 'fc-force-loop-ports' boot argument to force a particular topology.



## Corrective Action

(None).

## Syslog Message

Fibre Channel adapter port %s was forced to negotiate its topology as %s by topology-related boot arguments.

## Parameters

**port** (STRING): Name of the adapter port.

**topo** (STRING): Topology the adapter port was assigned.

# fci.config events

## fci.config.autoReboot

### Severity

NOTICE

### Description

This message occurs when the system reboots automatically to enable correct adapter operation after changing the state of a reconfigurable adapter.

### Corrective Action

No user action is required. The system is rebooting to enable the new adapter configuration.

### Syslog Message

Rebooting for Fibre Channel adapter %s to become available.

### Parameters

**adapters** (STRING): Affected adapters.

## fci.config.error

### Severity

ERROR

### Description

This message occurs when there is an error during the Fibre Channel adapter auto-reconfiguration process. The specified adapters will be unavailable until the problem is fixed.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Auto-reconfiguration for Fibre channel %s adapter(s) %s has failed.

### Parameters

**type** (STRING): Adapter type.

**adapters** (STRING): Affected adapter(s).

## **fci.config.missing**

### **Severity**

ERROR

### **Description**

This message occurs when the system cannot find Fibre Channel adapter configuration stored on disk, but the adapter is configured in the boot environment. The boot environment information will be used, but you should update the on-disk configuration to be consistent and to prevent the adapter from getting into an unexpected state if the boot environment settings are lost.

### **Corrective Action**

Use the 'fcadmin config' command to reapply the on-disk configuration to match the setting in the boot environment. If the on-disk configuration is not saved and the boot environment is lost, recovery of this adapter's port personality is not possible. No action is taken automatically to correct the problem.

### **Syslog Message**

Fibre Channel adapter %s is configured in the boot environment but the on-disk configuration information is missing.

### **Parameters**

**adapter** (STRING): Name of the affected Fibre Channel adapter.

## **fci.config.needsReboot**

### **Severity**

ERROR

### **Description**

This message occurs when the state of a reconfigurable adapter is changed, but the system must be rebooted to place the adapter into service.

### **Corrective Action**

Reboot the system to enable the new adapter configuration.

### **Syslog Message**

Reboot the system for Fibre channel %s adapter %s to become available.

### **Parameters**

**type** (STRING): Adapter type.

**adapters** (STRING): Affected adapter(s).

## **fci.config.offline**

### **Severity**

NOTICE

### **Description**

This message occurs when a configurable Fibre Channel adapter cannot be brought online because the configured adapter type needs to be changed.

## Corrective Action

Use the 'fcadmin config' utility to change the adapter configuration state and/or reboot the system to allow the PENDING adapter state change to take effect before bringing the adapter online.

## Syslog Message

Fibre Channel adapter %s is offline because it is in the %s state.

## Parameters

**adapter** (STRING): Name of the associated Fibre Channel adapter.

**state** (STRING): Configured adapter state.

## fci.config.state

### Severity

NOTICE

### Description

This message occurs when the state of a configurable Fibre Channel adapter changes.

### Corrective Action

(None).

## Syslog Message

Fibre Channel %s adapter %s is in the %s state.

## Parameters

**type** (STRING): Current adapter type.

**adapter** (STRING): Name of the associated Fibre Channel adapter.

**state** (STRING): Configured adapter state.

## fci.device events

### fci.device.general.reject

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Fibre Channel port transmits a frame that is rejected. The reject frame contains a reason code that can help diagnose the problem. In some cases, this might lead to a solution.

#### Corrective Action

(None).

#### Syslog Message

The frame transmitted by port %s was rejected by port\_id 0x%x with reason code 0x%x '%s', explanation code 0x%x '%s'.

#### Parameters

**port** (STRING): Name of the adapter port that transmitted the rejected frame.

**identifier** (INTEX): Destination port identifier of the rejected frame.

**reasonCode** (INTEX): Reason code byte provided for the exchange failure.  
**reasonString** (STRING): Reason code string provided for the exchange failure.  
**explanationCode** (INTEX): Explanation code byte provided for the exchange failure.  
**explanationString** (STRING): Explanation code string provided for the exchange failure.

## **fci.device.invalidate.soft.address**

### **Severity**

ERROR

### **Description**

This message occurs when a Fibre Channel disk drive takes a soft-assigned address and is being invalidated because soft-assigned addresses are not allowed.

### **Corrective Action**

Disk drives can take a soft-assigned address if the shelf IDs conflict. Make sure that all of the shelves in the loop have unique shelf IDs, and if different shelf types are being mixed, be sure to follow the documentation guidelines for mixing shelves on the same loop. After setting the shelves correctly, power-cycle the shelves so that the drives acquire the correct hard-assigned address.

### **Syslog Message**

Fibre Channel adapter %s is invalidating disk drive %s, which appears to have taken a soft address. Expected hard address %d (0x%x), assigned soft address %d (0x%x).

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.  
**deviceName** (STRING): Name of the associated Fibre Channel device.  
**hardLoopId** (INT): Hard loop ID that the device attempted to acquire.  
**hardALPA** (INTEX): Hard AL\_PA that the device attempted to acquire.  
**softLoopId** (INT): Soft loop ID that the device acquired.  
**softALPA** (INTEX): Soft AL\_PA that the device acquired.

## **fci.device.login.failure**

### **Severity**

ERROR

### **Description**

This message occurs when the Fibre Channel port login fails during the PLOGI or PRLI frame exchange. This can happen due to the reception of a Reject or Logout frame coming from the target, or due to a timeout.

### **Corrective Action**

Power-cycling the device might allow it to recover from this problem.

### **Syslog Message**

A login failure has occurred between Fibre Channel adapter %s and %s device %s, with error subcode %s.

### **Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.  
**deviceType** (STRING): Type of device.  
**deviceName** (STRING): Name of the device.

**code** (STRING): Error subcode provided by the firmware for the login failure, "N/A" if not provided.

## fci.device.login.reject

### Severity

INFORMATIONAL

### Description

This message occurs when the Fibre Channel port transmits a PLOGI frame that is rejected. The reject frame contains a reason code that can help diagnose the problem. In some cases, this might lead to a solution.

### Corrective Action

None.

### Syslog Message

The PLOGI transmitted by port %s was rejected by port\_id 0x%x with reason code 0x%x '%s', explanation code 0x%x '%s'.

### Parameters

**port** (STRING): Name of the adapter port that transmitted the PLOGI frame.

**identifier** (INHEX): Destination port identifier of the PLOGI frame.

**reasonCode** (INHEX): Reason code byte provided for the login failure.

**reasonString** (STRING): Reason code string provided for the login failure.

**explanationCode** (INHEX): Explanation code byte provided for the login failure.

**explanationString** (STRING): Explanation code string provided for the login failure.

## fci.device.loop.recovery

### Severity

ERROR

### Description

This message occurs when a device reports a break in the Fibre Channel loop, which could be due to a bad or unstable connection.

### Corrective Action

Check the cabling and verify that the disk drives and the shelf Fibre Channel modules on the loop are fully inserted. The problem is being caused by a connection in the loop upstream from the device.

### Syslog Message

Loop recovery event caused by the device upstream from %s device %s.

### Parameters

**deviceType** (STRING): Type of device.

**deviceName** (STRING): Name of the device.

## fci.discovery events

## **fci.discovery.timeout**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the Fibre Channel adapter driver cannot complete device discovery within the specified time. The driver allows ONTAP® software to continue booting and continues the discovery process in the background.

### **Corrective Action**

(None).

### **Syslog Message**

Device discovery timeout after %d seconds on adapter %s.

### **Parameters**

**time** (INT): Amount of time, in seconds, spent waiting for device discovery to be completed.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## **fci.fabric.events**

### **fci.fabric.max.ports**

#### **Severity**

ERROR

#### **Description**

This message occurs when a switch in the fabric attached to the Fibre Channel adapter exceeds the maximum number of ports allowed on a switch. Any devices on this switch will not be seen by the Fibre Channel adapter.

#### **Corrective Action**

Any devices on this switch will not be seen by the Fibre Channel adapter. If that is okay then this error can be ignored. If the devices need to be seen by the Fibre Channel adapter then they need to be connected to a switch that does not exceed the maximum number of ports supported per switch.

#### **Syslog Message**

Switch with WWN %s on Fibre Channel adapter %s exceeded the maximum number of ports supported: Maximum ports supported is %d and this switch has %d ports.

#### **Parameters**

**switchWWN** (STRING): World-Wide Name (WWN) of the switch that has exceeded the maximum number of supported ports.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**numMaxSwitches** (INT): Maximum number of ports supported per switch.

**numFabricSwitches** (INT): Number of ports on this switch.

### **fci.fabric.max.switches**

**Severity**

ERROR

**Description**

This message occurs when the fabric attached to the Fibre Channel adapter exceeds the maximum number of allowed switches. This configuration is not supported and will not work.

**Corrective Action**

Reduce the number of switches in the fabric.

**Syslog Message**

Invalid Fabric configuration detected on Fibre Channel adapter %s: Maximum switches supported is %d and this fabric contains %d switches.

**Parameters**

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**numMaxSwitches** (INT): Maximum number of switches supported in a fabric.

**numFabricSwitches** (INT): Number of switches in the fabric.

## fci.flash events

### fci.flash.preload.failed

**Severity**

ERROR

**Description**

This message occurs when there is a failure updating the Fibre Channel adapter flash preload table. If the preload table cannot be updated, the adapter might not have the best configuration settings for the PCI Express link.

**Corrective Action**

In most cases the adapter will still function normally. However, if you see PCI Express correctable errors on this adapter, you should replace the adapter.

**Syslog Message**

Failed to update flash preload table on Fibre Channel adapter %s with error %d.

**Parameters**

**port** (STRING): Name of the adapter port.

**errorCode** (INT): Update error code.

### fci.flash.preload.reboot

**Severity**

INFORMATIONAL

**Description**

This message occurs when the Fibre Channel adapter flash preload table has been updated and the controller is rebooting for the new settings to take effect.

**Corrective Action**

(None).

**Syslog Message**

Rebooting for updated Fibre Channel flash preload table to take effect.

**Parameters**

(None).

**fci.flash.preload.updated****Severity**

INFORMATIONAL

**Description**

This message occurs when the Fibre Channel adapter flash preload table is updated. The preload table contains configuration settings for the adapter.

**Corrective Action**

(None).

**Syslog Message**

Updated flash preload table on Fibre Channel adapter %s to version %s.

**Parameters**

**port** (STRING): Name of the adapter port.  
**version** (STRING): Preload table version.

**fci.fwupdate events****fci.fwupdate.AutoBoot****Severity**

NOTICE

**Description**

This message occurs when the Service Level Initiator Fibre Channel (SLIFC) Initiator driver successfully updates the firmware of the host bus adapter (HBA).

**Corrective Action**

None; the controller is rebooted automatically to enable the new HBA firmware.

**Syslog Message**

FC Initiator HBA %s firmware updated from version %s to %s. The adapter will be reset automatically to enable the new HBA firmware.

**Parameters**

**hba** (STRING): HBA where the firmware version update occurred.  
**currentversion** (STRING): Firmware version the HBA is running.  
**bundleversion** (STRING): Firmware version bundled with ONTAP® software.



## **fci.fwupdate.NeedBoot**

### **Severity**

NOTICE

### **Description**

This message occurs when the Service Level Interface Fibre Channel (SLIFC) Initiator driver successfully updates the firmware version of the host bus adapter (HBA), and requires a reboot of the controller to complete the process.

### **Corrective Action**

(None).

### **Syslog Message**

FC Initiator HBA %s firmware updated from version %s to %s. Reboot the controller to enable the new HBA firmware.

### **Parameters**

**hba** (STRING): HBA where the firmware update occurred.  
**currentversion** (STRING): Firmware version the HBA is running.  
**bundleversion** (STRING): Firmware version bundled with ONTAP® software.

## **fci.fwupdate.UpdDone**

### **Severity**

NOTICE

### **Description**

This message occurs when the Service Level Initiator Fibre Channel (SLIFC) Initiator driver successfully updates the firmware of the host bus adapter (HBA).

### **Corrective Action**

(None).

### **Syslog Message**

FC Initiator HBA %s firmware updated successfully from version %s to %s.

### **Parameters**

**hba** (STRING): HBA where the firmware update occurred.  
**currentversion** (STRING): Firmware version the HBA was running.  
**bundleversion** (STRING): Firmware version bundled with ONTAP® software.

## **fci.fwupdate.UpdErr**

### **Severity**

ERROR

### **Description**

This message occurs when the Service Level Initiator Fibre Channel (SLIFC) Initiator driver encounters an error while updating the firmware of the host bus adapter (HBA).

### Corrective Action

Contact NetApp technical support for further analysis.

### Syslog Message

FC Initiator HBA %s failed to update the firmware from version %s to %s with the status of %d.

### Parameters

**hba** (STRING): HBA where the firmware update failure occurred.  
**currentversion** (STRING): Firmware version the HBA is running.  
**bundleversion** (STRING): Firmware version bundled with ONTAP® software.  
**status** (INT): Status code of the firmware update process.

## fci.link events

### fci.link.break

#### Severity

ERROR

#### Description

This message occurs when the Fibre Channel adapter detects a break in the link.

#### Corrective Action

If the adapter is in use, check the cabling.

#### Syslog Message

Link break detected on Fibre Channel adapter %s.

#### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

### fci.link.break.shelf.fw.update

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Fibre Channel adapter detects a break in the link due to a shelf firmware update being performed.

#### Corrective Action

No corrective action is required. This is expected behavior when firmware is being downloaded to a shelf.

#### Syslog Message

Link break detected on Fibre Channel adapter %s while downloading shelf firmware.

#### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.link.error

### Severity

NOTICE

### Description

This message occurs when the Fibre Channel adapter cannot recover the link and the link is going offline.

### Corrective Action

(None).

### Syslog Message

Could not recover link on Fibre Channel adapter %s after %d seconds. Taking the adapter offline.

### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**downTime** (INT): Amount of time that the link is down.

## fci.loop.events

### fci.loop.stability

### Severity

ERROR

### Description

This message occurs when the Fibre Channel loop generates too many loop initialization events. This can be a sign of unstable hardware.

### Corrective Action

If the adapter is in use check cabling, and if it is connected to disk shelves check the seating of the LRC/ESH cards and disks. If this does not resolve the issue, contact NetApp technical support for assistance in diagnosing a Fibre Channel issue.

### Syslog Message

Fibre Channel adapter %s has a loop stability problem. %d loop-initialization events seen in %d seconds.

### Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**numLips** (INT): Number of loop-initialization events that occurred.

**numSeconds** (INT): Time interval that these loop-initialization events were seen.

## fci.mserr.events

### fci.mserr.general

### Severity

ERROR

## Description

This message occurs when the Fibre Channel fabric cannot successfully complete a request to the Management Server to retrieve configuration information. This query request is issued when the specified adapter is connected to a Fibre Channel fabric. Included in the event is the specific Management Server query along with text describing the reason and explanation codes returned by the fabric.

## Corrective Action

If the system happens to issue a Management Server query while the fabric is processing one or more configuration changes, the fabric might not successfully complete the request. After the fabric stabilizes, the Management Server can resume processing requests and the system automatically recovers from the error. This error is expected on McData switches that do not have the OSMS license installed, and devices on those switches cannot be seen by the Fibre Channel adapter. This error might also take place if the configuration has not been cabled properly, particularly if the explanation code translates into something like 'Invalid name for interconnect element or port'. If this error continues to occur after the fabric has stabilized and it is not happening on a McData switch without the OSMS license installed, check cabling, and contact NetApp technical support.

## Syslog Message

The Fibre Channel fabric Management Server request (%s) sent from adapter %s for WWN (%s) failed with reason code 0x%x (%s), explanation code 0x%x (%s), and response code 0x%x.

## Parameters

**msRequest** (STRING): Name of the failing Fibre Channel Management Server request.  
**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter (HBA).  
**queryWWN** (STRING): World Wide Name that was used for the query.  
**reasonCode** (INHEX): Reason code provided by the fabric for failure of the Management Server request.  
**reason** (STRING): Reason text provided by the fabric for failure of the Management Server request.  
**explanationCode** (INHEX): Explanation code provided by the fabric qualifying the stated reason.  
**explanation** (STRING): Explanation text provided by the fabric qualifying the stated reason.  
**responseCode** (INHEX): Response code provided by the fabric for failure of the specified Management Server request.

# fci.nserr events

## fci.nserr.general

### Severity

ERROR

### Description

This message occurs when the Fibre Channel fabric cannot successfully complete a Name Server query issued to obtain configuration and device information when an adapter is connected to the Fibre Channel fabric. Included in the event are the specific Name Server query and text describing the reason and explanation codes returned by the fabric.

### Corrective Action

If the system happens to issue a Name Server query while the fabric is processing one or more configuration changes, the fabric might not successfully complete the request. After the fabric stabilizes, the Name Server can resume processing requests and the system automatically recovers from the error. If this error continues to occur after the fabric has stabilized, contact NetApp technical support.

## Syslog Message

The Fibre Channel fabric Name Server request (%s) sent from adapter %s failed with reason code 0x%x (%s), explanation code 0x%x (%s) and response code 0x%x.

## Parameters

**nsRequest** (STRING): Name of the failing Fibre Channel Name Server request.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**reasonCode** (INHEX): Reason code provided by the fabric for failure of the Name Server request.

**reason** (STRING): Reason text provided by the fabric for failure of the Name Server request.

**explanationCode** (INHEX): Explanation code provided by the fabric qualifying the stated reason.

**explanation** (STRING): Explanation text provided by the fabric qualifying the stated reason.

**responseCode** (INHEX): Response code provided by the fabric for failure of the specified Name Server request.

## fci.nserr.noDevices

### Severity

ERROR

### Description

This message occurs when the fabric to which the specified adapter is attached reports that no Fibre Channel devices are present.

### Corrective Action

Check the fabric to ensure that all expected devices appear in the Name Server. If they appear in the Name Server but have not been reported to the specified adapter, review fabric zoning (if zoning is in use) to ensure that devices are present in the zone containing the adapter. If all expected devices are not present in the Name Server, verify the devices are powered-on, in an operational state, and switch connections/ports are operational.

## Syslog Message

The Fibre Channel fabric attached to adapter %s reports the presence of no Fibre Channel devices.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

## fci.nserr.peerWWN

### Severity

ERROR

### Description

This message occurs when the fabric connected to the specified adapter is zoned such that the adapter (or switch port if port zoning is in use) is not included in a fabric zone.

### Corrective Action

Check the fabric to ensure that the specified adapter is included in a fabric zone. If port zoning is in use, make sure that the switch port to which the adapter is attached is included in the active zone configuration. If WWN zoning is in use, ensure that the WWN of the specified adapter is included in the active zone configuration. Also make sure that the associated switch port is in an operational state.

## Syslog Message

Adapter %s is not in a fabric zone.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

# fci.set events

## fci.set.wwn

### Severity

INFORMATIONAL

### Description

This message occurs when the World-Wide Name (WWN) for this adapter is set to the value in the fci-wwn-"adapter\_name" bootarg.

### Corrective Action

(None).

## Syslog Message

Setting WWN for adapter %s to %s.

## Parameters

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter.

**wwn** (STRING): WWN that will be used by the adapter.

# fci.sfp events

## fci.sfp.invalid

### Severity

ERROR

### Description

This message occurs when an invalid Small Form-factor Pluggable (SFP) is inserted into the adapter. For example, a 10-GB SFP was inserted into an 8-GB FC adapter.

### Corrective Action

Replace the SFP with one that supports the proper speed.

## Syslog Message

The SFP in FC adapter %s is not valid. (%s)

## Parameters

**port** (STRING): Name of the adapter port.

**reason** (STRING): Reason for Invalid SFP.

# fci.wrong events

## fci.wrong.speed

### Severity

NOTICE

### Description

This message occurs when the fixed speed level requested by the fc-force-XG-ports bootarg is out of range for the corresponding FC port.

### Corrective Action

If another fixed speed is desired, adjust accordingly. Otherwise, the port will autonegotiate.

### Syslog Message

Fibre Channel adapter port %s does not support a speed setting of %u Gbps, which does not fall within port range %x.

### Parameters

**port** (STRING): Name of the adapter port.

**speed** (INT): Speed the adapter port was assigned.

**range** (INTHEX): Range of speeds the adapter port is capable of using.

# fcmon events

## fcmon.status events

### fcmon.status

#### Deprecated

Deprecated as of Data ONTAP 9.2 because FCMON process has been removed.

#### Severity

INFORMATIONAL

#### Description

This message occurs when FCMON starts up, halts or is requested to print it's current state.

#### Corrective Action

None.

#### Syslog Message

FCMON is %s

#### Parameters

**status** (STRING): The status of FCMON (initializing|running|sleeping|halted)



# fcp events

## fcp.service events

### fcp.service.shutdown

**Severity**

INFORMATIONAL

**Description**

The Fibre Channel Protocol (FCP) service has been shutdown.

**Corrective Action**

(None).

**Syslog Message**

FCP service shutdown

**Parameters**

(None).

### fcp.service.startup

**Severity**

INFORMATIONAL

**Description**

The Fibre Channel Protocol (FCP) service has been started.

**Corrective Action**

(None).

**Syslog Message**

FCP service startup

**Parameters**

(None).

# fcvi events

## fcvi.attr events

### fcvi.attr.changed

#### Severity

NOTICE

#### Description

This message occurs when any of the FCVI attributes changes.

#### Corrective Action

(None).

#### Syslog Message

FC-VI adapter: %s changed on port %s. %s

#### Parameters

**attr\_type** (STRING): Attribute type.

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

## fcvi.cm events

### fcvi.cm.timedout

#### Severity

NOTICE

#### Description

This message occurs when connection requests times out.

#### Corrective Action

(None).

#### Syslog Message

FC-VI adapter: Connection request timedout on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

#### Parameters

**port** (STRING): FCVI Interconnect NIC.

**qpname** (STRING): QP name.

**qpnum** (INT): QP index.

**system\_id** (LONGINT): Remote node's system id.

**info** (STRING): Additional info.

# fcvi.cross events

## fcvi.cross.connect.detected

### Severity

ERROR

### Description

This message occurs when the fcvi driver detects a cross connection

### Corrective Action

Check the FCVI interconnect cabling at the local and remote nodes. FCVI adapter ports fcvi\_device\_0s and fcvi\_device\_1s should not be connected to the same fabric. Contact NetApp technical support if the problem persists.

### Syslog Message

FC-VI adapter: Cross Connection detected between local port %s and remote port %s. Local System ID = %llu, Remote System ID = %llu.

### Parameters

**lname** (STRING): Local node's port name.  
**rname** (STRING): Remote node's port name.  
**lsid** (LONGINT): Local node's system id.  
**rsid** (LONGINT): Remote node's system id.

# fcvi.disabled events

## fcvi.disabled

### Severity

ALERT

### Description

This message occurs when the HA interconnect card is disabled by the ic-type environment variable.

### Corrective Action

Either remove the HA interconnect card from the controller or change the ic-type environment variable.

### Syslog Message

FC-VI adapter: Port %s has been disabled because of the ic-type environment variable.

### Parameters

**port** (STRING): FCVI Interconnect NIC.

# fcvi.fcrenabled events

## fcvi.fcrEnabled

**Severity**

NOTICE

**Description**

This message occurs if Fiber Channel Routing (FCR) functionality is enabled on the system.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Fiber Channel Routing (FCR) is enabled.

**Parameters**

(None).

## fcvi.initfailed events

### fcvi.initFailed

**Severity**

ERROR

**Description**

This message occurs when the driver fails to initialize a FC-VI port.

**Corrective Action**

Reboot the controller. Contact NetApp technical support if the problem persists.

**Syslog Message**

FC-VI adapter: %s failed for port %s.

**Parameters**

**operation** (STRING): Type of operation  
**port** (STRING): FCVI Interconnect NIC.

## fcvi.invalid events

### fcvi.invalid.advertisement

**Severity**

ERROR

**Description**

This message occurs when FCVI driver receives invalid advertisement.

**Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Port %s received invalid advertisement. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

## fcvi.invalid.sfp

### Severity

ERROR

### Description

This message occurs when the SFP (Small Form-Factor Pluggable) module of the Fibre Channel Virtual Interface (FC-VI) is invalid or incompatible.

### Corrective Action

Replace the SFP with a compatible one.

## Syslog Message

FC-VI adapter: Invalid SFP - Reason: %s on port %s.

## Parameters

**reason** (STRING): Reason for the incompatibility.

**port** (STRING): Interconnect NIC.

## fcvi.link events

### fcvi.link.down

### Severity

NOTICE

### Description

This message occurs when the physical link of the Fibre Channel Virtual Interface (FC-VI) adapter goes down.

### Corrective Action

1. Ensure that the FC-VI link has not been tampered with.
2. Verify that the physical status of the FC-VI adapter is "Up" by using the command "metrocluster interconnect adapter show".
3. If the configuration includes fabric switches, ensure that they are properly cabled and configured.

## Syslog Message

(None).

## Parameters

**alertId** (STRING): Alert ID.

**alertingResource** (STRING): Alerting resource unique name.

**alertingResourceName** (STRING): Alerting resource name.

**possibleEffect** (STRING): Possible effect for link transition.

**probableCause** (STRING): Probable cause of the event.

**probableCauseDescription** (STRING): Probable cause description of the event.  
**additionalInfo** (STRING): Any additional information for the event.

## fcvi.link.up

### Severity

NOTICE

### Description

This message occurs when the physical link of the Fibre Channel Virtual Interface (FC-VI) adapter comes up.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**alertId** (STRING): Alert ID.  
**alertingResource** (STRING): Alerting resource unique name.

## fcvi.qlgc events

### fcvi.qlgc.async.event.overflow

### Severity

ERROR

### Description

This message occurs when total number of async events are overflowed.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s async event overflowed. Resetting the adapter.

### Parameters

**port** (STRING): FCVI Interconnect NIC.

### fcvi.qlgc.connErr

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter detects a connection error. The FC-VI driver will attempt to retry the connection.

## Corrective Action

If this message repeats frequently, there may be an issue with the FC fabric or the adapter. Contact NetApp technical support.

## Syslog Message

FC-VI adapter: Port %s detects connection error. FC-VI opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**opcode\_str** (STRING): FCVI opcode.  
**opcode** (INHEX): FCVI opcode.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.  
**info** (STRING): Additional info.

## fcvi.qlgc.cpr.reset.cancelled

### Severity

ERROR

### Description

This message occurs when the CPR reset initiated by the FC-VI port got cancelled.

### Corrective Action

Reboot the filer. If issue still persists then replace the FCVI adapter.

### Syslog Message

FC-VI adapter: CPR reset initiated by port %s got cancelled. Offlining the port.

### Parameters

**port** (STRING): FCVI Interconnect adapter.

## fcvi.qlgc.cpr.reset.failed

### Severity

ERROR

### Description

This message occurs when the CPR reset initiated by the FC-VI port fails.

### Corrective Action

Reboot the filer. If issue still persists then replace the FCVI adapter.

### Syslog Message

FC-VI adapter: CPR reset initiated by port %s failed.

### Parameters

**port** (STRING): FCVI Interconnect adapter.

## **fcvi.qlgc.cpr.reset.unsupported**

### **Severity**

ERROR

### **Description**

This message occurs when the CPR reset is not supported by the adapter.

### **Corrective Action**

Reboot the filer. If issue still persists then replace the FCVI adapter.

### **Syslog Message**

FC-VI adapter: CPR reset is not supported by port %s.

### **Parameters**

**port** (STRING): FCVI Interconnect adapter.

## **fcvi.qlgc.FrameDiscardErr**

### **Severity**

ERROR

### **Description**

This message occurs when the FC-VI adapter firmware discards a received frame because of a receive frame error.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Discarded a received frame because of a receive frame error. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.

**qp\_name** (STRING): QP name.

**qp\_index** (INT): QP number.

**exchange\_id** (INTHEX): Exchange ID.

**address** (INTHEX): Address offset of the last good frame.

**system\_id** (LONGINT): Remote node's system id.

**info** (STRING): Additional info.

## **fcvi.qlgc.hard.reset.failed**

### **Severity**

ERROR

### **Description**

This message occurs when an FC-VI port's soft reset fails.



### Corrective Action

If CPR reset is not supported in the platform, then reboot the filer to recover from this issue. If issue persists then replace the FCVI adapter.

### Syslog Message

FC-VI adapter: Hard reset failed at port %s.

### Parameters

**port** (STRING): FCVI Interconnect adapter.

## fcvi.qlgc.invalid.data

### Severity

ERROR

### Description

This message occurs when FCVI driver receives invalid data from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid data. QP index = %d, completion status = %s(status code = %d), opcode = %s(%d), remote node's system id = %llu. %s

### Parameters

**port** (STRING): FCVI Interconnect NIC.

**qpnum** (INT): QP index.

**status** (STRING): Completion status.

**status\_code** (INT): Completion status code.

**opcode\_str** (STRING): FCVI opcode in string.

**opcode** (INT): FCVI opcode.

**systemID** (LONGINT): Remote node's system Id.

**info** (STRING): Additional info.

## fcvi.qlgc.invalid.exch.address

### Severity

ERROR

### Description

This message occurs when FCVI driver receives a response with invalid exchange address.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid exchange address for %s. QP name = %s, QP index = %d, expected exchange address = 0x%x, received exchange address = 0x%x. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**resp\_type** (STRING): Response type.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**exp\_oxid** (INTHEX): Expected exchange address.  
**rcvd\_oxid** (INTHEX): Received exchange address.  
**info** (STRING): Additional info.

## fcvi.qlgc.invalid.qp

### Severity

ERROR

### Description

This message occurs when FCVI driver receives invalid QP index from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid QP. Received QP index = %d, completion status = %s(status code = %d), entry type = %s(entry code = %d). %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qpnum** (INT): QP index.  
**status** (STRING): Completion status.  
**status\_code** (INT): Completion status code.  
**entry\_type** (STRING): Request/Response entry type.  
**entry\_code** (INT): Request/Response entry code.  
**info** (STRING): Additional info.

## fcvi.qlgc.invalid.response

### Severity

ERROR

### Description

This message occurs when FCVI driver receives invalid response from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid response. %s. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**error\_type** (STRING): Type of error.  
**info** (STRING): Additional info.

## fcvi.qlgc.ioErr

### Severity

ERROR

### Description

This message occurs when the FC-VI driver receives some IO error.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: FCVI driver on port %s received IO error. Status = %s(status code = 0x%x), FCVI opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu.

### Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code.  
**opcode\_str** (STRING): FCVI opcode in string.  
**opcode** (INHEX): FCVI opcode.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.

## fcvi.qlgc.login.timedout

### Severity

NOTICE

### Description

This message occurs when port login request times out.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Remote port login timedout on port %s. Remote port id = 0x%x, Remote port handle = %d.  
%s

### Parameters

**port** (STRING): FCVI Interconnect NIC.  
**portid** (INHEX): Remote port identifier.  
**porthandle** (INT): Remote port handle.  
**info** (STRING): Additional info.

## fcvi.qlgc.ls.reject

**Severity**

ERROR

**Description**

This message occurs when the FCVI firmware rejects a link service.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link service = %s(code = %d), rejected on port %s. %s

**Parameters**

**ls** (STRING): Link service name.

**ls\_code** (INT): Link service code.

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

**fcvi.qlgc.mbox.cmd.pending****Severity**

ERROR

**Description**

This message occurs when a mailbox command is pending for a long time.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Mailbox command '%s(code = 0x%x)' pending on port %s. %s

**Parameters**

**cmd** (STRING): Mailbox command.

**cmd\_code** (INTHEX): Mailbox command code.

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

**fcvi.qlgc.mbox.cmd.timedout****Severity**

ERROR

**Description**

This message occurs when a mailbox command times out.

**Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Mailbox command '%s(code = 0x%x)' timedout on port %s. %s

## Parameters

**cmd** (STRING): Mailbox command.  
**cmd\_code** (INTHEX): Mailbox command code.  
**port** (STRING): FCVI Interconnect NIC.  
**info** (STRING): Additional info.

## fcvi.qlgc.no.iu.buffer

### Severity

ERROR

### Description

This message occurs when the FCVI firmware identifies no more incoming IU buffers in the free pool.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Port %s doesn't have free IU buffer. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**info** (STRING): Additional info.

## fcvi.qlgc.no.recv.buffer

### Severity

ERROR

### Description

This message occurs when the FCVI firmware identifies no more receive buffers available for the incoming receive frames.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Port %s doesn't have free receive buffer. QP name = %s, QP index = %d, Message ID = 0x%x. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**msg\_id** (INTHEX): Message id of the incoming message.  
**info** (STRING): Additional info.

## fcvi.qlgc.oodErr

### Severity

ERROR

### Description

This message occurs when the FC-VI driver receives out of order completions.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: FCVI driver on port %s receives out of order completions. FCVI opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu.

### Parameters

**port** (STRING): Interconnect NIC.  
**opcode\_str** (STRING): FCVI opcode in string.  
**opcode** (INTHEX): FCVI opcode.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.

## fcvi.qlgc.PCIErr

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter firmware encounters a PCI error because it cannot write a frame to host memory.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: PCI error while writing a frame to host memory. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qp\_name** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **fcvi.qlgc.port.reset**

### **Severity**

ERROR

### **Description**

This message occurs when a FCVI port resets.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Resetting the port %s. %s

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.

**reason** (STRING): Reason for port reset.

## **fcvi.qlgc.qp.create.err**

### **Severity**

ERROR

### **Description**

This message occurs when QP creation fails.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Failed to create QP on port %s. Requested QP index = %d, completion status = %s(status code = %d). %s

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.

**qpnum** (INT): QP index.

**status** (STRING): Completion status.

**status\_code** (INT): Completion status code.

**info** (STRING): Additional info.

## **fcvi.qlgc.qp.destroy.failed**

### **Severity**

ERROR

### **Description**

This message occurs when the QP destroy request failed.

### **Corrective Action**

(None).

## Syslog Message

FC-VI adapter: QP destroy failed on port %s. QP index = %d. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.

**qpnum** (INT): QP index.

**info** (STRING): Additional info.

## fcvi.qlgc.qpErr

### Severity

ERROR

### Description

This message occurs when the FC-VI driver receives a QP error from the Firmware. The FC-VI driver will attempt to retry the connection.

### Corrective Action

If this message repeats frequently, there may be an issue with the FC fabric or the adapter. Contact NetApp technical support.

## Syslog Message

FC-VI adapter: FCVI driver on port %s received QP Error. QP name = %s, QP index = 0x%x, VI state = %s(state code = 0x%x), sub-error = %s(error code = 0x%x). Remote node's system id = %llu. This error has occurred %d times. %s

## Parameters

**port** (STRING): Interconnect Adapter.

**qpname** (STRING): QP name.

**qpnum** (INTHEX): QP number.

**state\_str** (STRING): QP state.

**state** (INTHEX): QP state code.

**suberror** (STRING): Sub-error.

**code** (INTHEX): Sub-error code.

**system\_id** (LONGINT): Remote node's system ID.

**errcnt** (INT): Number of times QP async error was encountered.

**info** (STRING): Additional info.

## fcvi.qlgc.rcvd.invalid.iu.buffer

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter receives invalid information unit buffer.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.



## Syslog Message

FC-VI adapter: Port %s received invalid IU buffer. Status = %s(status code = 0x%x), FCVI opcode = %s(0x%x), QP name = %s, QP index = %d, buffer handle = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INTEX): Completion status code.  
**opcode\_str** (STRING): FCVI opcode in string.  
**opcode** (INTEX): FCVI opcode.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**bufhdl** (INT): Buffer handle.  
**systemID** (LONGINT): Remote node's system ID.  
**info** (STRING): Additional info.

## fcvi.qlgc.receive.error

### Severity

ERROR

### Description

This message occurs when the FCVI firmware detects an error in the received frame.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Received bad frame on port %s. Mailbox-1 = %d, Mailbox-2 = %d. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**mbox1** (INT): Content of mailbox-1.  
**mbox2** (INT): Content of mailbox-2.  
**info** (STRING): Additional info.

## fcvi.qlgc.received.disconnect

### Severity

NOTICE

### Description

This message occurs when FCVI driver receives disconnect request from partner.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Disconnect request received on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## fcvi.qlgc.received.reject

### Severity

NOTICE

### Description

This message occurs when FCVI driver receives connection reject from partner.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Connect request rejected on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## fcvi.qlgc.recvIUErr

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter receives some error in received information unit.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: Port %s received IU error. Status = %s(status code = 0x%x), FCVI opcode = %s(0x%x), QP name = %s, QP index = %d, buffer handle = %d, Remote node's system id = %llu.

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code.  
**opcode\_str** (STRING): FCVI opcode in string.  
**opcode** (INHEX): FCVI opcode.

**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**bufhdl** (INT): Buffer handle.  
**systemID** (LONGINT): Remote node's system ID.

## **fcvi.qlgc.regDumped**

### **Severity**

NOTICE

### **Description**

This message occurs when the FCVI driver saves a diagnostic dump in a file.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Diagnostic dump for port %s saved to file %s.

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.  
**filename** (STRING): File to which the diagnostic dump was saved.

## **fcvi.qlgc.RespFrameErr**

### **Severity**

ERROR

### **Description**

This message occurs when the adapter firmware sends a response frame with the error bit set.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Firmware sent a response with the error bit set. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.  
**qpname** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **fcvi.qlgc.risc.paused**

### **Severity**

ERROR

**Description**

This message occurs when FCVI RISC is paused.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Port %s detected a RISC pause error. %s

**Parameters**

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

**fcvi.qlgc.rmt.link.down****Severity**

NOTICE

**Description**

This message occurs when FCVI driver lost link to partner port.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link to partner node over port %s is down. Partner port id = 0x%x, partner node's system id = %llu. %s

**Parameters**

**port** (STRING): FCVI Interconnect NIC.

**port\_id** (INHEX): Partner port id.

**system\_id** (LONGINT): Partner node's system id.

**info** (STRING): Additional info.

**fcvi.qlgc.rmt.link.up****Severity**

NOTICE

**Description**

This message occurs when FCVI driver gets link to partner port.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link to partner node over port %s is up. Partner port id = 0x%x, partner node's system id = %llu. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**port\_id** (INTHEX): Partner port id.  
**system\_id** (LONGINT): Partner node's system id.  
**info** (STRING): Additional info.

## fcvi.qlgc.sendIUError

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter fails to send an information unit.

### Corrective Action

No corrective action is required. The driver internally retries. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: Port %s sendIU status = %s(status code = 0x%x), FCVI opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INTHEX): Completion status code in hex.  
**opcode\_str** (STRING): FCVI opcode in string.  
**opcode** (INTHEX): FCVI opcode in hex.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## fcvi.qlgc.sent.disconnect

### Severity

NOTICE

### Description

This message occurs when FCVI driver sends disconnect request to remote partner.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Disconnect request sent on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.

**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **fcvi.qlgc.soft.reset.failed**

### **Severity**

ERROR

### **Description**

This message occurs when an FC-VI port's soft reset fails.

### **Corrective Action**

None.

### **Syslog Message**

FC-VI adapter: Soft reset failed at port %s.

### **Parameters**

**port** (STRING): FCVI Interconnect adapter.

## **fcvi.qlgc.system.error**

### **Severity**

ERROR

### **Description**

This message occurs when FCVI driver receives system error from FCVI firmware.

### **Corrective Action**

No corrective action is required. The driver internally resets the adapter. However, if you see this message frequently, there might be an issue with the adapter/firmware. Contact NetApp technical support.

### **Syslog Message**

FC-VI adapter: System Error on port %s. %s

### **Parameters**

**port** (STRING): FCVI Interconnect NIC.

**info** (STRING): Additional info.

## **fcvi.qlgc.TimeOutErr**

### **Severity**

ERROR

### **Description**

This message occurs when the adapter firmware cannot send a request.

### **Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Firmware could not send a request. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**qp\_name** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## fcvi.qlgc.transfer.error

### Severity

ERROR

### Description

This message occurs when the FCVI firmware detects a transfer error.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: %s transfer error on port %s. Affected queue number = %d. %s

## Parameters

**trans\_type** (STRING): Transfer type.  
**port** (STRING): FCVI Interconnect NIC.  
**queue\_num** (INT): Queue number.  
**info** (STRING): Additional info.

## fcvi.qlgc\_regdumpfailed events

### fcvi.qlgc\_regDumpFailed

### Severity

NOTICE

### Description

This message occurs when the HA interconnect driver could not save a diagnostic dump.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Failed to save diagnostic dump for port %s. %s.

## Parameters

**port** (STRING): FCVI Interconnect NIC.  
**reason** (STRING): Reason for the failure.

# fcvi.unsupported events

## fcvi.unsupported.config

### Severity

ERROR

### Description

This message occurs when a network adapter port behavior type is set to use FC-VI on a platform that has not been validated for its use and is therefore unsupported.

### Corrective Action

Change the port behavior type to a value other than "FC-VI". Use the "system hardware unified-connect show" command to view network adaptor port information, including the current behavior type. Use the "system hardware unified-connect modify" command to change the network adapter port behavior type.

### Syslog Message

FC-VI port behavior type not supported on network adapter port %s.

## Parameters

**port** (STRING): The network adapter port whose behavior type is set to use FC-VI.

# fcvi.update events

## fcvi.update.link.state

### Severity

NOTICE

### Description

This message occurs when there is a change in physical link state of the Fibre Channel Virtual Interface (FC-VI) adapter.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Physical link is %s on port %s.

## Parameters

**link\_state** (STRING): New link state.  
**port** (STRING): Interconnect NIC.



# ffm events

## ffm.aggr events

### ffm.aggr.attach.done

#### Severity

NOTICE

#### Description

This message occurs when an aggregate attach operation is complete. The event displays the service ID of the attached Firefly Service and the generation number, which is the number of times this aggr has been unattached.

#### Corrective Action

(None).

#### Syslog Message

Aggregate ID: %u has been attached to a Firefly service ID: %u, the current attach generation of the aggregate is %u.

#### Parameters

**aggrID** (INT): Unique identifier of the aggregate.

**fireflyID** (INT): Unique identifier of the Firefly Service the aggregate is attached to.

**attachGeneration** (INT): Current attach generation of the aggregate.

### ffm.aggr.attach.started

#### Severity

NOTICE

#### Description

This message occurs when an aggregate attach operation is started. The event displays the service ID of the attached Firefly Service and the generation number, which is the number of times this aggr has been unattached.

#### Corrective Action

(None).

#### Syslog Message

Attachment of an aggregate ID: %u to a Firefly service ID: %u has been started, the current attach generation of the aggregate is %u.

#### Parameters

**aggrID** (INT): Unique identifier of the aggregate.

**fireflyID** (INT): Unique identifier of the Firefly Service the aggregate is being attached to.

**attachGeneration** (INT): Current attach generation of the aggregate.

## ffm.aggr.detach.done

### Severity

NOTICE

### Description

This message occurs when an aggregate detach operation is completed. The event displays the service ID of the attached Firefly Service and the generation number, which is the number of times this aggr has been unattached.

### Corrective Action

(None).

### Syslog Message

Aggregate ID: %u has been detached from a Firefly service ID: %u, the current attach generation of the aggregate is %u.

### Parameters

**aggrID** (INT): Unique identifier of the aggregate.

**fireflyID** (INT): Unique identifier of the Firefly Service the aggregate is detached from.

**attachGeneration** (INT): Current attach generation count of the aggregate.

## ffm.aggr.detach.started

### Severity

NOTICE

### Description

This message occurs when an aggregate detach operation is started. The event displays the service ID of the attached Firefly Service and the generation number, which is the number of times this aggr has been unattached.

### Corrective Action

(None).

### Syslog Message

Detachment of an aggregate ID: %u from a Firefly service ID: %u has been started, the current attach generation of the aggregate is %u.

### Parameters

**aggrID** (INT): Unique identifier of the aggregate.

**fireflyID** (INT): Unique identifier of the Firefly Service the aggregate is being detached from.

**attachGeneration** (INT): Current attach generation of the aggregate.

## ffm.attach events

### ffm.attach.outage.end

### Severity

NOTICE

## Description

This message occurs when volumes on an aggregate become available to the protocol clients near the end of the attach phase of volume relocation.

## Corrective Action

(None).

## Syslog Message

Protocol outage for volumes on aggregate "%s" has ended at %llu, %llu ms after the attach callback (including %llu ms for replay).

## Parameters

**aggr\_name** (STRING): Name of the aggregate.

**sk\_msecs** (LONGINT): Timestamp of this event.

**attach\_protocol\_outage\_duration** (LONGINT): Duration in milliseconds from the moment of the attach callback to this event.

**attach\_replay\_duration** (LONGINT): Duration of the nvlog replay before the protocol outage ends.

## ffm.attach.outage.summary

### Severity

NOTICE

### Description

This message occurs when the attach phase of graceful or ungraceful volume relocation finishes.

### Corrective Action

(None).

### Syslog Message

Attach of volumes on aggregate "%s" has started at %llu and ended at %llu. Total attach time: %llu ms.

### Parameters

**aggr\_name** (STRING): Name of the aggregate.

**attach\_start\_sk\_msecs** (LONGINT): Timestamp when the attach starts.

**attach\_end\_sk\_msecs** (LONGINT): Timestamp of this event.

**attach\_op\_total** (LONGINT): Duration in milliseconds from the moment of the attach callback to this event.

## ffm.dataip events

### ffm.dataip.failedToHost

#### Severity

NOTICE

#### Description

This message occurs when a data IP address was failed to host on the node. Data access through this IP address is not available.

### Corrective Action

(None).

### Syslog Message

Data IP address %s was failed to host by the NFS protocol service on the node %s (ID: %u).

### Parameters

**dataip** (STRING): The IP address that was failed to host.

**nodename** (STRING): Name of the node on which the data IP was failed to host.

**nodeid** (INT): Unique identifier of the node on which the data IP was failed to host.

## ffm.dataip.notUnhosted

### Severity

NOTICE

### Description

This message occurs when a data IP address is not unhosted on the node.

### Corrective Action

(None).

### Syslog Message

The NFS protocol service could not unhost data IP address "%s" on node "%s" (ID: %u).

### Parameters

**dataip** (STRING): IP address that could not be unhosted.

**nodename** (STRING): Name of the node on which the data IP could not be unhosted.

**nodeid** (INT): Unique identifier of the node on which the data IP could not be unhosted.

## ffm.dataip.successfullyHosted

### Severity

NOTICE

### Description

This message occurs when a data IP address was successfully hosted on the node. Data access through this IP address is now available.

### Corrective Action

(None).

### Syslog Message

Data IP address %s was hosted by the NFS protocol service on the node %s (ID: %u).

### Parameters

**dataip** (STRING): The IP address that was hosted.

**nodename** (STRING): Name of the node on which the data IP was hosted.

**nodeid** (INT): Unique identifier of the node on which the data IP was hosted.

## ffm.dataip.unhostSuccess

### Severity

NOTICE

### Description

This message occurs when a data IP address is successfully unhosted from the node. Data access through this IP address is not available.

### Corrective Action

(None).

### Syslog Message

The NFS protocol service unhosted data IP address "%s" on node "%s" (ID: %u).

### Parameters

**dataip** (STRING): IP address that was unhosted.

**nodename** (STRING): Name of the node on which the data IP was unhosted.

**nodeid** (INT): Unique identifier of the node on which the data IP was unhosted.

## ffm.detach events

### ffm.detach.outage.start

#### Severity

NOTICE

#### Description

This message occurs when a volume becomes unavailable to protocol traffic during a graceful volume relocation.

#### Corrective Action

(None).

#### Syslog Message

Protocol outage of volume "%s" on aggregate "%s" has started. Timestamp: %llu, %llu ms after the detach callback.

#### Parameters

**volume\_name** (STRING): Name of the volume.

**aggr\_name** (STRING): Name of the aggregate.

**sk\_msecs** (LONGINT): Timestamp of this event.

**detach\_lead\_time** (LONGINT): Duration in milliseconds from the moment of the detach callback to this event.

### ffm.detach.outage.summary

#### Severity

NOTICE

## Description

This message occurs when the detach phase of graceful volume relocation finishes.

## Corrective Action

(None).

## Syslog Message

Detach protocol outage has completed for volumes on aggregate %s. Timestamp: %llu, estimated protocol outage duration: %llu ms, total detach duration: %llu ms.

## Parameters

**aggr\_name** (STRING): Name of the aggregate.

**detach\_end\_sk\_msecs** (LONGINT): Timestamp of this event.

**est\_detach\_protocol\_outage\_duration** (LONGINT): Estimated duration in milliseconds from the start to the end of the protocol outage for the detach operation.

**detach\_op\_total** (LONGINT): Duration in milliseconds from the moment of the detach callback to this event.

# ffm.init events

## ffm.init.firefly

### Severity

NOTICE

### Description

This message occurs when the Firefly manager is initiated during Firefly boot.

### Corrective Action

(None).

### Syslog Message

Firefly manager initiated.

### Parameters

(None).

# ffm.restart events

## ffm.restart.firefly

### Severity

NOTICE

### Description

This message occurs when the Firefly manager restarts Firefly after detecting abnormal behavior. This restart avoids a core dump, which would not help address the abnormal behavior. When Firefly comes back up, aggregates will be re-attached.

**Corrective Action**

(None).

**Syslog Message**

Firefly manager restart: %s

**Parameters**

**reason** (STRING): Reason to restart.

# ffrespmon events

## ffrespmon.aggr events

### ffRespMon.aggr.offline

#### Severity

NOTICE

#### Description

This message occurs when an aggregate is detected to be offline due to an unresponsive Firefly Service on a node.

#### Corrective Action

(None).

#### Syslog Message

Aggregate (ID: %u) attached (attaching: %u) to a Firefly Service (ID: %u), running on node %s (ID: %u, Uuid: %s), is offline.

#### Parameters

**aggrID** (INT): Unique identifier of the aggregate.

**attaching** (INT): Boolean value to denote whether the Firefly Service was in the process of attaching the aggregate.

**fireflyID** (INT): Unique identifier of the Firefly Service to which the aggregate is attached.

**nodeName** (STRING): Name of the node hosting the Firefly Service.

**nodeID** (INT): Unique identifier of the node the hosting the Firefly Service.

**nodeUuid** (STRING): Universal unique identifier of the node the hosting the Firefly Service.

## ffrespmon.nfs events

### ffRespMon.nfs.unresponsive

#### Severity

NOTICE

#### Description

This message occurs when a data interface's IP address is unreachable. The NFS service might be affected and unresponsive.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s (ID: %u), hosted on node %s (ID: %u, Uuid: %s), is currently not reachable.

#### Parameters

**dataip** (STRING): IP address of the data interface.

**dataipID** (INT): Unique identifier of the data IP.

**nodeName** (STRING): Name of the node hosting given data IP.



**nodeID** (INT): Unique identifier of the node hosting given data IP.

**nodeUuid** (STRING): Universal unique identifier of the node hosting given data IP.

# fg events

## fg.admin events

### fg.admin.wait.long

#### Severity

NOTICE

#### Description

This message occurs when the data component takes longer than 15 seconds to update its view of the FlexGroup volume after an admin operation.

#### Corrective Action

(None).

#### Syslog Message

FlexGroup volume "%s" (MSID "%lu") in Vserver "%s" took %lu seconds to update the "%s" property during the "%s" operation.

#### Parameters

**volumeName** (STRING): Name of the FlexGroup volume.

**volumeMsid** (LONGINT): Master Data Set ID (MSID) of the FlexGroup volume.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup volume resides.

**numSeconds** (LONGINT): Number of seconds that the data component took to update.

**property** (STRING): Volume property being updated.

**operation** (STRING): Admin operation performed before the data component update.

## fg.aggr events

### fg.aggr.mixed.types.error

#### Severity

ERROR

#### Description

This message occurs when the constituents of a FlexGroup volume are located on different types of aggregates (FabricPool-enabled aggregates and non-FabricPool-enabled aggregates) for more than one week. It is reported once per week if a FlexGroup volume remains in this state. Leaving a FlexGroup volume in a mixed-aggregate state might lead to unbalanced data placement and unexpected performance issues.

#### Corrective Action

Move all the constituents of the FlexGroup volume to the same aggregate type, i.e., either FabricPool-enabled aggregates or non-FabricPool-enabled aggregates, using the 'volume move start' command with the 'allow-mixed-aggr-types' option set to true.

#### Syslog Message

The constituents of FlexGroup volume '%s' with UUID '%s' have been located on different types of aggregates (FabricPool-enabled aggregates and non-FabricPool-enabled aggregates) for more than one

week.

## Parameters

**vol** (STRING): Name of the FlexGroup volume.  
**uuid** (STRING): UUID of the FlexGroup volume.

## fg.aggr.mixed.types.notice

### Severity

NOTICE

### Description

This message occurs when the constituents of a FlexGroup volume are located on different types of aggregates (FabricPool-enabled aggregates and non-FabricPool-enabled aggregates). It is reported once per hour. Leaving a FlexGroup volume in a mixed-aggregate state might lead to unbalanced data placement and unexpected performance issues.

### Corrective Action

Move all the constituents of the FlexGroup volume to the same aggregate type, i.e., either FabricPool-enabled aggregates or non-FabricPool-enabled aggregates, using the 'volume move start' command with the 'allow-mixed-aggr-types' option set to true.

### Syslog Message

The constituents of FlexGroup volume '%s' with UUID '%s' are located on different types of aggregates (FabricPool-enabled aggregates and non-FabricPool-enabled aggregates).

## Parameters

**vol** (STRING): Name of the FlexGroup volume.  
**uuid** (STRING): UUID of the FlexGroup volume.

## fg.convert events

### fg.convert.critical.fail

### Severity

ALERT

### Description

This message occurs when the conversion of a FlexVol® volume into a FlexGroup fails. The FlexVol volume is inoperable, and can no longer serve data.

### Corrective Action

(Call support).

### Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu failed. Reason: %s. The FlexVol can no longer serve data. Reason: %s.

## Parameters

**volumeName** (STRING): Name of the volume being converted to a FlexGroup.  
**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

**reason** (STRING): Reason the conversion failed.

**rollbackFailedReason** (STRING): Reason the conversion could not roll back after failing to convert.

## fg.convert.fail

### Severity

ERROR

### Description

This message occurs when the conversion of a FlexVol® volume into a FlexGroup fails. The FlexVol volume is still operable.

### Corrective Action

Ensure that the FlexVol volume is a valid candidate for the convert operation, and then retry the command. If the error persists, contact NetApp technical support.

### Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu failed. Reason: %s

### Parameters

**volumeName** (STRING): Name of volume being converted to a FlexGroup.

**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

**reason** (STRING): Reason why the conversion failed.

## fg.convert.ksmf.timeout

### Severity

NOTICE

### Description

This message occurs when the operation to convert a FlexVol® volume into a FlexGroup times out when updating the volume's RAID label. The system retries the conversion automatically.

### Corrective Action

(None).

### Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu has timed out while updating the volume's RAID label.

### Parameters

**volumeName** (STRING): Name of the volume being converted to a FlexGroup.

**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

## fg.convert.start

### Severity

NOTICE

### Description

This message occurs when the conversion of a FlexVol® volume to a FlexGroup begins.

## Corrective Action

(None).

## Syslog Message

Starting conversion of FlexVol volume %s to new FlexGroup with MSID %llu, UUID %s.

## Parameters

**volumeName** (STRING): Name of volume being converted to a FlexGroup.

**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

**flexgroupUuid** (STRING): UUID of the new FlexGroup.

## fg.convert.success

### Severity

NOTICE

### Description

This message occurs when the conversion of a FlexVol® volume into a FlexGroup succeeds.

## Corrective Action

(None).

## Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu successful.

## Parameters

**volumeName** (STRING): Name of volume being converted to a FlexGroup.

**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

## fg.create events

### fg.create.workflow.non.auto

### Severity

NOTICE

### Description

This message occurs when a user creates a FlexGroup volume that has a workflow with a type that is not "auto".

## Corrective Action

(None).

## Syslog Message

FlexGroup volume "%s" in Vserver "%s" was created with workflow type "%s"

## Parameters

**volumeName** (STRING): Name of the FlexGroup volume.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup volume resides.

**workflow** (STRING): The type of workflow for the FlexGroup volume.

# fg.inactive events

## fg.inactive.dc.only.delete

### Severity

NOTICE

### Description

This message occurs when the system automatically deletes an inactive FlexGroup constituent that existed only in the data component.

### Corrective Action

(None).

### Syslog Message

Constituent with DISD "%lu" of FlexGroup volume "%s" (MSID "%lu") in Vserver "%s", which exists only in the data component and is inactive, has been automatically deleted.

### Parameters

**constituentDsid** (LONGINT): Data Set ID (DSID) of the deleted constituent.

**volumeName** (STRING): Name of the FlexGroup volume owning the deleted constituent.

**volumeMsid** (LONGINT): Master Data Set ID (MSID) of the FlexGroup volume owning the deleted constituent.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup volume resides.

## fg.inactive.delete

### Severity

NOTICE

### Description

This message occurs when the system automatically deletes an inactive FlexGroup constituent.

### Corrective Action

(None).

### Syslog Message

Constituent "%lu" (DISD "%lu") of FlexGroup volume "%s" (MSID "%lu") in Vserver "%s" is inactive and has been automatically deleted.

### Parameters

**flexgroupIndex** (LONGINT): FlexGroup index of the deleted constituent.

**constituentDsid** (LONGINT): Data Set ID (DSID) of the deleted constituent.

**volumeName** (STRING): Name of the FlexGroup volume owning the deleted constituent.

**volumeMsid** (LONGINT): Master Data Set ID (MSID) of the FlexGroup volume owning the deleted constituent.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup volume resides.

# fg.inodes events

## fg.inodes.member.allOK

### Severity

NOTICE

### Description

This message occurs when conditions that led to previous "fg.inodes.member.nearlyFull" and "fg.inodes.member.full" events no longer apply for any constituent in this FlexGroup. All constituents within this FlexGroup have sufficient inodes for normal operation.

### Corrective Action

(None).

### Syslog Message

All constituents in FlexGroup %s (fg-uuid %s) have sufficient inodes for normal operation.

### Parameters

**object\_name** (STRING): Human-readable name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

## fg.inodes.member.full

### Severity

ALERT

### Description

This message occurs when a constituent with a FlexGroup has run out of inodes. New files cannot be created on this constituent. This might lead to an overall imbalanced distribution of content across the FlexGroup.

### Corrective Action

Adding capacity to the FlexGroup by using the "volume modify -files +X" command is the best way to solve this problem. Alternatively, deleting files from the FlexGroup may work, but it is difficult to determine which files have landed on which constituent.

### Syslog Message

Constituent %llu in FlexGroup %s (fg-uuid %s) is out of inodes.

### Parameters

**constituent\_dsid** (LONGINT): DSID of the constituent that is out of inodes.

**object\_name** (STRING): Human-readable name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

## fg.inodes.member.nearlyFull

### Severity

ALERT

## Description

This message occurs when a constituent within a FlexGroup is almost out of inodes. This constituent will receive far fewer new create requests than average, which might impact the FlexGroup's overall performance, because those requests are routed to constituents with more inodes.

## Corrective Action

Adding capacity to the FlexGroup by using the "volume modify -files +X" command is the best way to solve this problem. Alternatively, deleting files from the FlexGroup might work, although it can be difficult to determine which files have landed on which constituent.

## Syslog Message

Constituent %llu in FlexGroup %s (fg-uuid %s) is nearly out of inodes.

## Parameters

**constituent\_dsid** (LONGINT): DSID of the constituent that is almost out of inodes.

**object\_name** (STRING): Human-readable name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

# fg.member events

## fg.member.below.threshold

### Severity

ALERT

### Description

This message occurs when a members available space is below the threshold. ONTAP will be reporting zero space for the entire flexgroup.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact NetApp technical support for further assistance.

### Syslog Message

FlexGroup volume "%s" (Size: "%lu", Usage: "%lu"%%) has member volume "%s" (Size: "%lu", Usage: "%lu"%%) that is at or below the free space threshold of "%lu" blocks.

### Parameters

**flexgroupName** (STRING): Name of the FlexGroup volume.

**flexgroupSize** (LONGINT): Size in blocks of the FlexGroup volume.

**flexgroupUsage** (LONGINT): Usage percentage of the FlexGroup volume.

**memberName** (STRING): Name of the FlexGroup member volume.

**memberSize** (LONGINT): Size in blocks of the FlexGroup member volume.

**memberUsage** (LONGINT): Usage percentage of the FlexGroup member volume.

**thresholdSize** (LONGINT): Size of threshold that has been reached.

## fg.member.elastic.sizing

### Severity

NOTICE



## Description

This message occurs when a constituent volume in a FlexGroup volume undergoes elastic sizing either to restore balance among constituent volumes or to resize them to accommodate space needs.

## Corrective Action

(None).

## Syslog Message

Constituent volume "%u" in FlexGroup volume "%s" was resized from %lu bytes to %lu bytes. This constituent volume (%s) %s by resizing the constituent volume "%s" (dsid "%u"). This constituent volume is %llu bytes in size with %llu bytes used (%u percent full). The FlexGroup volume is %llu bytes in size with %llu bytes used (%u percent full).

## Parameters

**constituent\_dsid** (INT): DSID of the constituent volume undergoing elastic sizing.  
**object\_name** (STRING): Name of the FlexGroup volume.  
**current\_size** (LONGINT): Original block size of the constituent volume.  
**new\_size** (LONGINT): New block size of the constituent volume.  
**constituent\_name** (STRING): Name of the constituent volume.  
**operation** (STRING): Indicates whether the constituent volume is being grown or shrunk.  
**partner\_name** (STRING): Name of the constituent volume's partner.  
**partner\_dsid** (INT): DSID of the constituent volume's partner.  
**size** (LONGINT): Size in bytes of the constituent volume.  
**used** (LONGINT): Size in bytes of the used space of the constituent volume.  
**percent\_full** (INT): Percentage of filled space in the constituent volume.  
**flexgroup\_size** (LONGINT): Size in bytes of the FlexGroup volume.  
**flexgroup\_used** (LONGINT): Size in bytes of the used space in the FlexGroup volume.  
**flexgroup\_percent\_full** (INT): Percentage of filled space in the FlexGroup volume.  
**event\_time** (LONGINT): When the event occurred, shown in system time.

# fg.modify events

## fg.modify.workflow

### Severity

NOTICE

### Description

This message occurs when a user modifies the workflow type of a FlexGroup volume.

### Corrective Action

(None).

### Syslog Message

The workflow type for FlexGroup volume "%s" in Vserver "%s" was modified from "%s" to "%s".

### Parameters

**volumeName** (STRING): Name of the FlexGroup volume.  
**vserverName** (STRING): Name of the Vserver on which the FlexGroup volume resides.  
**oldWorkflow** (STRING): The old workflow type of the FlexGroup volume.  
**newWorkflow** (STRING): The new workflow type for the FlexGroup volume.

# fg.move events

## fg.move.unmirrored.aggr

### Severity

NOTICE

### Description

This message occurs when a user moves a FlexGroup constituent to an unmirrored aggregate on nodes in a MetroCluster(tm) configuration. However, unmirrored aggregates, as a rule, do not have synchronous disaster protection.

### Corrective Action

(None).

### Syslog Message

FlexGroup constituent %s was moved to unmirrored aggregate %s on Vserver %s.

### Parameters

**volumeName** (STRING): Name of the FlexGroup constituent moved to the unmirrored aggregate.

**aggregateName** (STRING): Name of the unmirrored aggregate.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup constituent resides.

# fg.multidelete events

## fg.multidelete.inval.snapid

### Severity

NOTICE

### Description

This message occurs when a Snapshot copy in an invalid state is deleted from a FlexGroup volume. For example, if a Snapshot copy creation fails on some constituent volumes of a FlexGroup volume, but succeeds on others, it is deemed invalid and deleted by ONTAP.

### Corrective Action

(None).

### Syslog Message

Invalid Snapshot copy ID "%u" on FlexGroup volume "%s" is being deleted (UUID: %s, MSID: %llx).

### Parameters

**snapid** (INT): ID of the invalid Snapshot copy to be deleted.

**flexgroup\_name** (STRING): Name of the FlexGroup volume.

**flexgroup\_uuid** (STRING): UUID of the FlexGroup volume.

**flexgroup\_msid** (LONGINHEX): Master Data Set ID (MSID) of the FlexGroup volume.

# fg.plex events

## fg.plex.del.unmirror.aggr

### Severity

NOTICE

### Description

This message occurs when a user unmirrors an aggregate (plex delete) that contains a FlexGroup constituent on nodes in a MetroCluster(tm) configuration. However, unmirrored aggregates, as a rule, do not have synchronous disaster protection.

### Corrective Action

(None).

### Syslog Message

Aggregate %s, which contains a FlexGroup constituent, was unmirrored.

### Parameters

**aggregateName** (STRING): Name of the aggregate that is unmirrored.

## fg.provision events

### fg.provision.unmirrored.aggr

### Severity

NOTICE

### Description

This message occurs when a user provisions a FlexGroup on unmirrored aggregates on nodes in a MetroCluster(tm) configuration. However, unmirrored aggregates, as a rule, do not have synchronous disaster protection.

### Corrective Action

(None).

### Syslog Message

FlexGroup %s was provisioned on unmirrored aggregates on Vserver %s.

### Parameters

**volumeName** (STRING): Name of the FlexGroup created on the unmirrored aggregates.

**vserverName** (STRING): Name of the Vserver on which the FlexGroup resides.

## fg.quota events

### fg.quota.cannot.enforce

### Severity

NOTICE

### Description

This message occurs when a FlexGroup cannot enforce its quota as expected.

## Corrective Action

(None).

## Syslog Message

FlexGroup %s (fg-uuid %s) cannot enforce quota due to %s.

## Parameters

**object\_name** (STRING): Name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

**reason** (STRING): Reason for lack of quota enforcement.

## fg.quota.enforcement.on

### Severity

NOTICE

### Description

This message occurs when a file operation discovers that the quota enforcement state of a FlexGroup has changed to "on".

### Corrective Action

(None).

### Syslog Message

FlexGroup %s (fg-uuid %s) quota enforcement state is on.

### Parameters

**object\_name** (STRING): Name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

## fg.quota.on.failed

### Severity

ERROR

### Description

This message occurs when enabling quotas (quota on an operation) on a FlexGroup constituent fails.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact technical support.

### Syslog Message

Failed to execute the quota on operation on FlexGroup constituent %s in Vserver %s. Reason: %s.

### Parameters

**volumeName** (STRING): Name of the constituent.

**vserverName** (STRING): Name of the Vserver.

**reason** (STRING): Reason for the failure.

## fg.quota.usermapping.result

### Severity

NOTICE

### Description

This message occurs when the quota mapper decides whether to map the Windows quota record and the UNIX quota record of a user into a single multiuser record.

### Corrective Action

(None).

### Syslog Message

%s (uuid %s) tid %d: %s.

### Parameters

**constituent\_name** (STRING): Name of the FlexGroup constituent.

**object\_uuid** (STRING): UUID of the FlexGroup constituent.

**qtree** (LONGINT): Qtree identifier (tid).

**result** (STRING): User mapping result.

## fg.recover events

### fg.recover.fail

#### Severity

ERROR

#### Description

This message occurs when the recovery of a FlexGroup volume back into a FlexVol® fails. The FlexGroup volume is still operable.

#### Corrective Action

Ensure that the FlexGroup volume is a valid candidate for the recovery operation, and then retry the command. If the error persists, contact NetApp technical support.

#### Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu failed. Reason: %s

#### Parameters

**volumeName** (STRING): Name of volume being recovered back into a FlexVol.

**flexgroupMsid** (LONGINTEX): MSID of the recovered FlexVol.

**reason** (STRING): Reason why the recovery failed.

### fg.recover.start

#### Severity

NOTICE

## Description

This message occurs when the recovery of a FlexGroup volume back into a FlexVol® begins.

## Corrective Action

(None).

## Syslog Message

Starting recovery of FlexGroup volume %s back into a FlexVol® with MSID %llu.

## Parameters

**volumeName** (STRING): Name of volume being recovered back into a FlexVol.

**volMsid** (LONGINTEX): MSID of the recovered FlexVol.

## fg.recover.success

### Severity

NOTICE

### Description

This message occurs when the recovery of a FlexGroup volume back into a FlexVol® succeeds.

### Corrective Action

(None).

### Syslog Message

Conversion of FlexVol volume %s to FlexGroup with MSID %llu successful.

### Parameters

**volumeName** (STRING): Name of volume being recovered back into a FlexVol.

**volMsid** (LONGINTEX): MSID of the recovered FlexVol.

## fg.space events

### fg.space.member.allOK

#### Severity

NOTICE

#### Description

This message occurs when conditions that led to previous "fg.space.member.nearlyFull" and "fg.space.member.full" events no longer apply. All constituents within this FlexGroup have sufficient space for normal operation.

#### Corrective Action

(None).

#### Syslog Message

The FlexGroup %s (fg-uuid %s) no longer has any full constituents.

## Parameters

**object\_name** (STRING): Human-readable name of the FlexGroup.

**object\_uuid** (STRING): UUID of the FlexGroup.

## fg.space.member.full

### Severity

ALERT

### Description

This message occurs when a constituent volume within a FlexGroup volume is now full. Although new files can still be created within the FlexGroup volume and some files can still be expanded, none of the files that are stored on this constituent volume can be modified. The user might see random out-of-space errors as a consequence when trying to perform write operations on this FlexGroup volume.

### Corrective Action

Adding capacity to the FlexGroup volume by using the "volume modify -size +X" command is the best way to solve this problem. Alternatively, deleting files from the FlexGroup volume might work, but it is difficult to determine which files have landed on which constituent volume.

### Syslog Message

Constituent volume "%llu" in FlexGroup volume "%s" (fg-uuid "%s") is out of space. This constituent volume ("%s") is %llu bytes in size with %llu bytes used (%u percent full). The FlexGroup volume is %llu bytes in size with %llu bytes used (%u percent full).

## Parameters

**constituent\_dsid** (LONGINT): DSID of the constituent volume that is out of space.

**object\_name** (STRING): Human-readable name of the FlexGroup volume.

**object\_uuid** (STRING): UUID of the FlexGroup volume.

**constituent\_name** (STRING): Name of the constituent volume.

**size** (LONGINT): Size in bytes of the constituent volume.

**used** (LONGINT): Size in bytes of the used space of the constituent volume.

**percent\_full** (INT): Percentage of filled space in the constituent volume.

**flexgroup\_size** (LONGINT): Size in bytes of the FlexGroup volume.

**flexgroup\_used** (LONGINT): Size used in bytes of the FlexGroup volume.

**flexgroup\_percent\_full** (INT): Percentage of filled space in the FlexGroup volume.

**event\_time** (LONGINT): When the event occurred, shown in system time.

## fg.space.member.nearlyFull

### Severity

ALERT

### Description

This message occurs when a constituent volume within a FlexGroup volume is nearly out of space. Although files can be created and expanded, if the constituent volume runs out of space, the user's ability to append to and modify those files that reside on this constituent volume will be impeded.

### Corrective Action

Adding capacity to the FlexGroup volume by using the "volume modify -size +X" command is the best way to solve this problem. Alternatively, deleting files from the FlexGroup volume might work, but it is difficult to determine which files have landed on which constituent volume.

## Syslog Message

Constituent volume "%llu" in FlexGroup volume "%s" (fg-uuid "%s") is nearly out of space. This constituent volume ("%s") is %llu bytes in size with %llu bytes used (%u percent full). The FlexGroup volume is %llu bytes in size with %llu bytes used (%u percent full).

## Parameters

**constituent\_dsid** (LONGINT): DSID of the constituent volume that is nearly out of space.  
**object\_name** (STRING): Human-readable name of the FlexGroup volume.  
**object\_uuid** (STRING): UUID of the FlexGroup volume.  
**constituent\_name** (STRING): Name of the constituent volume.  
**size** (LONGINT): Size in bytes of the constituent volume.  
**used** (LONGINT): Size in bytes of the used space of the constituent volume.  
**percent\_full** (INT): Percentage of filled space in the constituent volume.  
**flexgroup\_size** (LONGINT): Size in bytes of the FlexGroup volume.  
**flexgroup\_used** (LONGINT): Size used in bytes of the FlexGroup volume.  
**flexgroup\_percent\_full** (INT): Percentage of filled space in the FlexGroup volume.  
**event\_time** (LONGINT): When the event occurred, shown in system time.

# fg.statechange events

## fg.stateChange

### Severity

NOTICE

### Description

This message occurs when the FlexGroup has transitioned from one state to another. This message might not occur immediately after the state transition occurred.

### Corrective Action

(None).

### Syslog Message

FlexGroup %s (fg-uuid %s) has changed state from %s to %s.

### Parameters

**object\_name** (STRING): Human-readable name of the FlexGroup.  
**object\_uuid** (STRING): UUID of the FlexGroup.  
**old\_state** (STRING): State that the FlexGroup was in before the EMS event.  
**new\_state** (STRING): State that the FlexGroup was in at the time of the EMS event.

# fg.svmdr events

## fg.svmdr.qos.wklds.insuff

### Severity

NOTICE

### Description

This message occurs during FlexGroup volume replication in an SVM DR setup when the available QoS workloads are insufficient on the secondary cluster.



### Corrective Action

Contact NetApp technical support.

### Syslog Message

Available QoS workloads are insufficient to support replicated FlexGroup "%s" in Vserver "%s".

### Parameters

**volume** (STRING): FlexGroup name.

**vserver** (STRING): Vserver name.

## fg.unconfig events

### fg.unconfig.consts.exist

#### Severity

ERROR

#### Description

This message occurs when a FlexGroup detects constituents that are not properly configured and have not been removed after a sufficient timeout. This usually means that a failed attempt to expand the volume also failed to delete constituents.

#### Corrective Action

Contact NetApp technical support for assistance.

#### Syslog Message

FlexGroup %s has had constituents that are not properly configured for longer than %d seconds.

#### Parameters

**object\_name** (STRING): Name of the FlexGroup.

**seconds** (INT): Timeout period in seconds.

## fg.unconvert events

### fg.unconvert.critical.fail

#### Severity

ALERT

#### Description

This message occurs when a member of a FlexGroup fails to be converted into an independent FlexVol® volume. The original state of this member could not be restored, so the FlexGroup itself is inoperable, and requires repair.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

FlexGroup member #%u of %u in FlexGroup %s failed to be converted into an independent FlexVol volume. The FlexGroup is inoperable. Reason: %s

## Parameters

**fgindex** (INT): FlexGroup index of the volume being converted to a FlexVol volume.

**fgmembers** (INT): Total number of members in this FlexGroup.

**fgname** (STRING): Name of the FlexGroup containing the member being converted to a FlexVol volume.

**reason** (STRING): Reason why the conversion failed.

## fg.unconvert.fail

### Severity

ERROR

### Description

This message occurs when a member of a FlexGroup has fails to be converted into an independent FlexVol® volume. The FlexGroup is still functional.

### Corrective Action

Ensure that the FlexGroup is a valid candidate for the unconvert operation, and then retry the command. If the error persists, contact NetApp technical support.

### Syslog Message

FlexGroup member #%u of %u in FlexGroup %s failed to be converted into an independent FlexVol volume.  
Reason: %s

## Parameters

**fgindex** (INT): FlexGroup index of the volume being converted to a FlexVol volume.

**fgmembers** (INT): Total number of members in this FlexGroup.

**fgname** (STRING): Name of the FlexGroup containing the member being converted to a FlexVol volume.

**reason** (STRING): Reason why the conversion failed.

## fg.unconvert.start

### Severity

NOTICE

### Description

This message occurs when a member of a FlexGroup, usually the root constituent, is being converted to an independent FlexVol® volume.

### Corrective Action

(None).

### Syslog Message

Starting conversion of FlexGroup member #%u of %u in FlexGroup %s into an independent FlexVol volume.

## Parameters

**fgindex** (INT): FlexGroup index of the volume being converted to a FlexVol volume.

**fgmembers** (INT): Total number of members in this FlexGroup.

**fgname** (STRING): Name of the FlexGroup containing the member being converted to a FlexVol volume.

## fg.unconvert.success

### Severity

NOTICE

### Description

This message occurs when a member of a FlexGroup is successfully converted to an independent FlexVol® volume.

### Corrective Action

(None).

### Syslog Message

FlexGroup member #%u of %u in FlexGroup %s was successfully converted into an independent FlexVol volume.

### Parameters

**fgindex** (INT): FlexGroup index of the volume being converted to a FlexVol volume.

**fgmembers** (INT): Total number of members in this FlexGroup.

**fgname** (STRING): Name of the FlexGroup containing the member being converted to a FlexVol volume.

## fg.writeback events

### fg.writeback.delta.checker

### Severity

NOTICE

### Description

This message occurs when an inconsistency is discovered during a writeback. The writeback is ignored, so that the inconsistency does not get transferred to the cache origin.

### Corrective Action

(None).

### Syslog Message

Writeback operation failed for %s at line %d (%s).

### Parameters

**fh** (STRING): File handle of the affected inode.

**source** (INT): Location of the failing writeback validation test.

**state** (STRING): Description of the inconsistency.

# fgr events

## fgr.metadata events

### fgr.metadata.repaired

#### Severity

NOTICE

#### Description

This message occurs when the system detects an inconsistency in the volume capacity rebalancing metadata. As a consequence, the inconsistent metadata is deleted and new metadata will be repopulated according to latest volume capacity rebalancing state.

#### Corrective Action

None.

#### Syslog Message

The system detected an inconsistency in the volume capacity rebalancing metadata associated with volume %s:%s.

#### Parameters

**vserveruuid** (STRING): Universal unique identifier (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

# fips events

## fips.post events

### fips.post.failure

#### Severity

ERROR

#### Description

This message occurs when an unrecoverable error prevents initialization of the cryptographic subsystem into FIPS mode. No SSL communication will be possible with the storage system in this state. Other ONTAP features may also be affected by the lack of cryptographic operations.

#### Corrective Action

To operate in non FIPS cryptographic mode, turn off FIPS mode for the SSL interface using command "security config modify -interface SSL -is-fips-enabled false" and restart the system. For more information or assistance, contact NetApp technical support.

#### Syslog Message

SSL interface disabled due to failure of FIPS mode cryptographic services.

#### Parameters

(None).

# firmware events

## firmware.file events

### firmware.file.invalid

#### Severity

ERROR

#### Description

This message occurs when the system detects an invalid disk or shelf firmware file. This file cannot be used to update firmware on any disk or shelf device. This message repeats every 24 hours as long as the invalid file is present. There might be other invalid files on the system, and messages for those files are suppressed if they are detected within 24 hours of this event.

#### Corrective Action

Check your support provider for more information about obtaining the latest disk and shelf firmware. If the problem persists, contact NetApp technical support.

#### Syslog Message

%s %s failed verification due to: %s.

#### Parameters

**fwFileType** (STRING): Type of the firmware file, for example disk or shelf firmware file.

**fwFileName** (STRING): Name of the firmware file.

**reason** (STRING): Reason for mismatch.

# flexcache events

## flexcache.attrib events

### flexcache.attrib.sync.failed

#### Severity

ALERT

#### Description

This message occurs when when syncing properties between an origin of a FlexCache volume and its FlexCache volume fails, and due to which FlexCache volume may not be able to proceed with accurate data.

#### Corrective Action

Use the "volume flexcache sync-properties" command on the cluster hosting the FlexCache volume to sync the origin volume properties with the FlexCache volume.

#### Syslog Message

Failed to sync the FlexCache configuration for cluster %s in Vserver %s hosting FlexCache volume %s.

#### Parameters

**cacheCluster** (STRING): Name of the cluster hosting the FlexCache volume.  
**cacheVserver** (STRING): Name of the Vserver hosting the FlexCache volume.  
**cacheVolName** (STRING): Name of FlexCache volume for which sync failed.

### flexcache.attrib.sync.passed

#### Severity

NOTICE

#### Description

This message occurs when origin of a FlexCache volume properties sync succeeds for all connected FlexCache volumes.

#### Corrective Action

(None).

#### Syslog Message

Sync of origin of a FlexCache volume %s in Vserver %s succeeded for all connected FlexCache volumes.

#### Parameters

**originVolName** (STRING): Origin of a FlexCache volume.  
**originVserver** (STRING): Name of the Vserver hosting the origin of a FlexCache volume.

### flexcache.attrib.sync.warn

#### Severity

NOTICE

## Description

This message occurs when syncing properties between an origin of a FlexCache volume and its FlexCache volume fails, and the sync operation will be retried.

## Corrective Action

Use the "volume flexcache sync-properties" command on the cluster hosting the FlexCache volume to sync the origin volume properties with the FlexCache volume.

## Syslog Message

Failed to sync the FlexCache configuration for cluster %s in Vserver %s hosting FlexCache volume %s. Operation will be retried.

## Parameters

**cacheCluster** (STRING): Name of the cluster hosting the FlexCache volume.

**cacheVserver** (STRING): Name of the Vserver hosting the FlexCache volume.

**cacheVolName** (STRING): Name of FlexCache volume for which sync failed.

# flexcache.autogrow events

## flexcache.autogrow.fail

### Severity

ERROR

### Description

This message occurs when automatically growing a FlexCache volume fails because the total cluster-wide FlexCache provisioned space has exceeded the licensed limit.

### Corrective Action

To increase the licensed capacity, install a new license in the cluster.

### Syslog Message

Autosize of FlexCache volume '%s%s%s' failed because the total cluster-wide FlexCache provisioned space has exceeded the licensed limit in the cluster %s.

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**cluster\_uuid** (STRING): UUID of the Cluster.

# flexcache.cacheconnected events

## flexcache.cacheConnected

### Severity

INFORMATIONAL



## Description

This message occurs when a FlexCache® origin volume establishes communication with its cache volume.

## Corrective Action

No action required.

## Syslog Message

FlexCache origin volume with Name %s MSID %lu on Vserver Name %s UUID %s is connected with FlexCache volume with Vserver UUID %s and MSID %lu.

## Parameters

**originVolumeName** (STRING): Origin Volume name.  
**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.  
**originVserverName** (STRING): Vserver name of the origin volume.  
**originVserverUuid** (STRING): Vserver UUID of the origin volume.  
**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.  
**cacheMsid** (LONGINT): MSID of the FlexCache volume.

# flexcache.cachedisconnected events

## flexcache.cacheDisconnected

### Severity

EMERGENCY

### Description

This message occurs when a FlexCache® origin volume cannot communicate with its cache volume. The cache volume might therefore present stale data to its clients until the connection is reestablished.

### Corrective Action

Verify that the origin volume can connect with the FlexCache volume. If both volumes are within the same cluster, check the health of the owning nodes by using the "system node show" command. If the volumes are on different clusters, ensure that the intercluster LIFs on each cluster can connect by using the "network ping" command.

### Syslog Message

FlexCache origin volume with Name %s MSID %lu on Vserver Name %s UUID %s is disconnected from FlexCache volume with Vserver UUID %s and MSID %lu.

### Parameters

**originVolumeName** (STRING): Origin Volume name.  
**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.  
**originVserverName** (STRING): Vserver name of the origin volume.  
**originVserverUuid** (STRING): Vserver UUID of the origin volume.  
**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.  
**cacheMsid** (LONGINT): MSID of the FlexCache volume.

# flexcache.capacity events

## flexcache.capacity.exceeded

### Severity

ERROR

### Description

This message occurs when the total cluster-wide FlexCache provisioned space has exceeded the licensed limit.

### Corrective Action

To increase the licensed capacity, install a new license in the cluster.

### Syslog Message

The total cluster-wide FlexCache provisioned space has exceeded the licensed limit. Cluster id: %s, Current usage: %s, Licensed capacity: %s.

### Parameters

**cluster\_uuid** (STRING): UUID of the Cluster.

**used\_capacity** (STRING): Total FlexCache provisioned space in the cluster.

**licensed\_capacity** (STRING): FlexCache licensed capacity.

## flexcache.conf events

### flexcache.conf.refreshFailed

### Severity

ALERT

### Description

This message occurs when FlexCache configuration refresh fails. If this event happens, FlexCache will not be able to reach to its origin which will cause FlexCache to serve only cached data. This message occurs at most once every 30 minutes.

### Corrective Action

Use the "volume flexcache config-refresh" command on the cluster hosting the FlexCache volume to refresh the origin volume configuration on FlexCache volume.

### Syslog Message

Failed to refresh FlexCache configuration for peer vserver UUID %s and peer volume MSID %lu.  
isPeerOrigin : %u.

### Parameters

**peerVserverUuid** (STRING): Peer vserver UUID in string format for which refresh failed.

**peerVolMsid** (LONGINT): Peer volume MSID for which refresh failed.

**isPeerOrigin** (INT): Origin peer volume or cache peer volume for which refresh failed.

## flexcache.config events

## flexcache.config.sync.failed

### Severity

ERROR

### Description

This message occurs when the FlexCache volume expand operation fails to synchronize the FlexCache configuration information between the newly created constituent volume and the origin cluster.

### Corrective Action

Use the "volume flexcache config-refresh" command on the origin cluster to update the FlexCache configuration.

### Syslog Message

Volume expand failed to create the cache configuration for FlexCache volume "%s" in Vserver "%s" on the origin of the FlexCache cluster. Error: %s.

### Parameters

**cacheVolName** (STRING): Name of the FlexCache volume.

**cacheVserver** (STRING): Name of the Vserver hosting the FlexCache volume.

**errorString** (STRING): Error string.

## flexcache.encryption events

### flexcache.encryption.mismatch.with.origin

### Severity

ALERT

### Description

This message occurs when a plaintext FlexCache volume has been created for an encrypted origin volume.

### Corrective Action

(None).

### Syslog Message

A plaintext FlexCache volume "%s" in Vserver "%s" has been created for origin volume "%s" in Vserver "%s".

### Parameters

**originVolumeName** (STRING): Origin Volume Name.

**cacheVolumeName** (STRING): Cache Volume Name.

**cacheVserverName** (STRING): Cache Vserver Name.

**originVserverName** (STRING): Origin Vserver Name.

## flexcache.featuredisabled events

### flexcache.featureDisabled

**Severity**

ERROR

**Description**

This message occurs when an attempt is made to mount a FlexCache volume. FlexCache is no longer supported, please remove any remaining FlexCache volumes left over from upgrading the cluster.

**Corrective Action**

Delete all FlexCache volumes in the cluster. They can be displayed by using the "volume flexcache show" command.

**Syslog Message**

(None).

**Parameters**

**volume** (STRING): FlexCache volume involved.

## flexcache.mcc events

### flexcache.mcc.refresh.failed

**Severity**

ERROR

**Description**

This message occurs when a FlexCache® configuration refresh operation fails during MetroCluster(tm) switchover and switchback. Data access through FlexCache volumes might be impacted as a result.

**Corrective Action**

First, verify that the two Vservers in question are peered by using the "vserver peer show" command. Next, enter the "flexcache config-refresh -peer-server <peer-vserver> -peer-volume <peer-volume> -peer -endpoint-type <peer-endpoint-type>" command on a node within the cluster that directly manages the peer endpoint.

**Syslog Message**

FlexCache configuration refresh failed on peer Vserver %s for peer volume %s and peer endpoint %s on cluster %s.

**Parameters**

**peeredVserverUuid** (STRING): UUID of the peered Vserver.

**peerVolumeName** (STRING): Name of the peer volume.

**peerEndPointType** (STRING): Peer endpoint type.

**clusterUuid** (STRING): UUID of the cluster with the peer endpoint.

**localVserverUuid** (STRING): UUID of the local Vserver.

## flexcache.name events

### flexcache.name.sync.failed

## Severity

ERROR

## Description

This message occurs when the FlexCache® volume rename operation fails to synchronize the FlexCache name information to the origin cluster. Because of this failure, the "volume flexcache origin show-caches" command shows the old FlexCache volume name for the origin relationship.

## Corrective Action

Use the "volume flexcache sync-properties" command on the FlexCache cluster to update the FlexCache name on the origin cluster. Make sure that the FlexCache cluster and the origin of the FlexCache cluster are running on the same ONTAP® build.

## Syslog Message

Volume rename failed to update the new FlexCache volume "%s" on Vserver "%s" to the origin of the FlexCache cluster. Error: %s.

## Parameters

**cacheVolName** (STRING): Name of the FlexCache volume.

**cacheVserver** (STRING): Name of the Vserver hosting the FlexCache volume.

**errorString** (STRING): Error string.

# flexcache.nearly events

## flexcache.nearly.full

### Severity

ERROR

### Description

This message occurs when the total cluster-wide FlexCache provisioned space has nearly reached the licensed limit.

### Corrective Action

To increase the licensed capacity, install a new license in the cluster.

### Syslog Message

The total cluster-wide FlexCache provisioned space has nearly reached the licensed limit. Cluster id: %s, Current usage: %s, Licensed capacity: %s.

### Parameters

**cluster\_uuid** (STRING): UUID of the Cluster.

**used\_capacity** (STRING): Total FlexCache provisioned space in the cluster.

**licensed\_capacity** (STRING): FlexCache licensed capacity.

# flexcache.origin events

## flexcache.origin.config.sync.failed

**Severity**

ERROR

**Description**

This message occurs when the origin of the FlexCache volume expand operation fails to synchronize the configuration information between the newly created constituent volume and the cache cluster.

**Corrective Action**

Use the "volume flexcache config-refresh" command on the cache cluster to update the origin of the FlexCache configuration.

**Syslog Message**

Volume expand failed to create the origin configuration for the origin of the FlexCache volume "%s" in Vserver "%s" on the FlexCache cluster. Error: %s.

**Parameters**

**originVolName** (STRING): Name of the origin of the FlexCache volume.

**cacheVserver** (STRING): Name of the Vserver hosting the origin of the FlexCache volume.

**errorString** (STRING): Error string.

**flexcache.origin.size.sync.failed.disconnected.mode****Severity**

ERROR

**Description**

This message occurs when syncing the size property of the origin volume between an origin of a FlexCache volume and its FlexCache volume fails because the cluster containing the FlexCache volume is not reachable.

**Corrective Action**

Make sure that the cluster containing the FlexCache volume is reachable. Make sure that the FlexCache cluster and the origin of the FlexCache cluster are running on the same ONTAP® build.

**Syslog Message**

Failed to sync origin size property for FlexCache Cluster "%s". Error: %s

**Parameters**

**originClusterUuid** (STRING): Cache Cluster UUID.

**errorString** (STRING): Error string.

**flexcache.origin.size.sync.failed.for.cache.cluster****Severity**

ERROR

**Description**

This message occurs when syncing the size property of the origin volume between an origin of a FlexCache volume and its FlexCache volume fails because we encountered some error on the cluster containing the FlexCache volume.

### Corrective Action

Look for an EMS on the cache cluster and take the corrective action needed. Make sure that the FlexCache cluster and the origin of the FlexCache cluster are running on the same ONTAP® build.

### Syslog Message

Failed to sync origin size property for Cache Cluster "%s".

### Parameters

**cacheClusterUuid** (STRING): Cache Cluster UUID.

## flexcache.origin.size.sync.failed.for.cache.volume

### Severity

ERROR

### Description

This message occurs when syncing the size property of the origin volume between an origin of a FlexCache volume and its FlexCache volume fails, and the sync operation will be retried.

### Corrective Action

Take the appropriate action depending on the error message. Make sure that the FlexCache cluster and the origin of the FlexCache cluster are running on the same ONTAP® build.

### Syslog Message

Failed to sync origin size property for FlexCache volume "%s" in Vserver "%s". Error: %s.

### Parameters

**volumeName** (STRING): Cache Volume Name.

**vserverName** (STRING): Cache Vserver Name.

**errorString** (STRING): Error string.

## flexcache.origincleanup events

### flexcache.origincleanup.fail

### Severity

ERROR

### Description

This message occurs when the origin of a FlexCache® volume cleanup job fails to clean up the origin volume. All future FlexCache operations on the origin volume will fail.

### Corrective Action

Use the "volume flexcache origin cleanup-cache-relationship -force-retry true" command on the origin cluster to cleanup origin of a FlexCache volume. If the problem persists then contact technical support for assistance.

### Syslog Message

Origin volume cleanup job failed to clean up the origin volume of the FlexCache volume %s in Vserver %s.

## Parameters

**originVolName** (STRING): Name of the origin volume of the FlexCache volume.  
**originVserver** (STRING): Name of the Vserver hosting the origin of a FlexCache volume.

## flexcache.origincleanup.retry

### Severity

ERROR

### Description

This message occurs when the RAL cleanup at the origin volume of a FlexCache® volume fails due to timeout. All future FlexCache operations on the origin volume will fail.

### Corrective Action

Use the "volume flexcache origin cleanup-cache-relationship -force-retry true" command on the origin cluster to cleanup origin of a FlexCache volume. If the problem persists then contact technical support for assistance.

### Syslog Message

RAL cleanup incomplete for the origin volume of the FlexCache volume %s in Vserver %s.

## Parameters

**originVolName** (STRING): Name of the origin volume of the FlexCache volume.  
**originVserver** (STRING): Name of the Vserver hosting the origin of a FlexCache volume.

## flexcache.originconnected events

### flexcache.originConnected

### Severity

INFORMATIONAL

### Description

This message occurs when a FlexCache® volume establishes communication with its origin volume.

### Corrective Action

No action required.

### Syslog Message

FlexCache volume with Name %s MSID %lu on Vserver Name %s UUID %s is connected with its origin volume with Vserver UUID %s and MSID %lu.

## Parameters

**cacheVolumeName** (STRING): Volume Name of the FlexCache volume.  
**cacheMsid** (LONGINT): MSID of the FlexCache volume.  
**cacheVserverName** (STRING): Vserver Name of the FlexCache volume.  
**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.  
**originVserverUuid** (STRING): Vserver UUID of the origin volume.  
**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.



# flexcache.originDisconnected events

## flexcache.originDisconnected

### Severity

EMERGENCY

### Description

This message occurs when a FlexCache® volume cannot communicate with its origin volume. The cache volume might therefore present stale data to its clients. Client I/O to the cache volume might also not appear to respond if the client is requesting data that has not been cached.

### Corrective Action

Verify that the FlexCache volume can connect with its origin volume. If both volumes are within the same cluster, check the health of the owning nodes by using the "system node show" command. If the volumes are on different clusters, ensure that the intercluster LIFs on each cluster can connect by using the "network ping" command.

### Syslog Message

FlexCache volume with Name %s MSID %lu on Vserver Name %s UUID %s is disconnected from its origin volume with Vserver UUID %s and MSID %lu.

### Parameters

**cacheVolumeName** (STRING): Volume Name of the FlexCache volume.

**cacheMsid** (LONGINT): MSID of the FlexCache volume.

**cacheVserverName** (STRING): Vserver Name of the FlexCache volume.

**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.

**originVserverUuid** (STRING): Vserver UUID of the origin volume.

**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.

# flexcache.originfull events

## flexcache.originFull

### Severity

ALERT

### Description

This message occurs when a FlexCache® origin volume cannot serve data to cache volume because origin volume is full. The cache volume might therefore not be able to serve data to its clients until the origin has free space.

### Corrective Action

Ensure that the origin volume has free space.

### Syslog Message

FlexCache origin volume with Name %s MSID %lu on Vserver Name %s UUID %s may not be able to serve data to FlexCache volume with Vserver UUID %s and MSID %lu because origin volume is full.

## Parameters

**originVolumeName** (STRING): Origin Volume name.  
**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.  
**originVserverName** (STRING): Vserver name of the origin volume.  
**originVserverUuid** (STRING): Vserver UUID of the origin volume.  
**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.  
**cacheMsid** (LONGINT): MSID of the FlexCache volume.

# flexcache.originfull\_1 events

## flexcache.originFull\_1

### Severity

ALERT

### Description

This message occurs when a FlexCache® volume cannot retrieve data from its origin volume because origin volume is full. Client I/O to the cache volume might not appear to respond if the client is requesting data that has not been cached.

### Corrective Action

Ensure that the origin volume has free space.

### Syslog Message

FlexCache volume with Name %s MSID %lu on Vserver Name %s UUID %s may not be able to retrieve data from its origin volume with Vserver UUID %s and MSID %lu because origin volume is full.

## Parameters

**cacheVolumeName** (STRING): Volume Name of the FlexCache volume.  
**cacheMsid** (LONGINT): MSID of the FlexCache volume.  
**cacheVserverName** (STRING): Vserver Name of the FlexCache volume.  
**cacheVserverUuid** (STRING): Vserver UUID of the FlexCache volume.  
**originVserverUuid** (STRING): Vserver UUID of the origin volume.  
**originMsid** (LONGINT): Master data set identifier (MSID) of the origin volume.

# flexcache.softevict events

## flexcache.softEvict.failed

### Severity

ERROR

### Description

This message occurs when soft evicting a cache volume fails. The cache volume might therefore present stale data to its clients.

### Corrective Action

(None).

**Syslog Message**

Soft eviction of constituent DSID %lu failed.

**Parameters**

**constituentDsid** (LONGINT): Constituent's DSID.

**flexcache.softEvict.success****Severity**

INFORMATIONAL

**Description**

This message occurs on successfully soft evicting the cache volume.

**Corrective Action**

No Action Required.

**Syslog Message**

Constituent DSID %lu is successfully soft evicted.

**Parameters**

**constituentDsid** (LONGINT): Constituent's DSID.

# fli events

## fli.lun events

### fli.lun.import.checkpoint

#### Severity

ERROR

#### Description

This message occurs when a checkpoint record delete fails. Checkpoint is metadata information stored with the NetApp LUN attribute stream for tracking import progress. An attempt is made to delete this metadata during import stop or delete operations.

#### Corrective Action

This record can remain in the LUN attribute stream without any harm. It could occupy ~100-200 bytes of space in the LUN attribute stream. If that is an issue, a support personnel could be contacted to delete this record.

#### Syslog Message

Import checkpoint record could not be deleted for LUN %s - Reason:%s

#### Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**FailureReason** (STRING): Internal error reason.

### fli.lun.import.complete

#### Severity

INFORMATIONAL

#### Description

This message occurs when a foreign LUN import is complete.

#### Corrective Action

(None).

#### Syslog Message

Import of foreign LUN %s of size %lld bytes from array model %s belonging to vendor %s is successfully completed. Destination NetApp LUN is %s.

#### Parameters

**DiskSerialNumber** (STRING): Foreign disk serial number.

**Size** (LONGINT): Size of the foreign LUN.

**ForeignArrayModelID** (STRING): Foreign array model ID.

**ForeignArrayVendor** (STRING): Foreign array vendor name.

**LUNSerialNumber** (STRING): Serial number of the LUN.

## fli.lun.import.failed

### Severity

ERROR

### Description

This message occurs when a foreign LUN import goes into failed state. 1. Data could not be read from the source(foreign LUN). 2. Data could not be written to the destination(NetApp LUN). 3. This is an Internal/metadata related error.

### Corrective Action

1. Verify the connectivity to the foreign disk.
2. Verify that the destination volume is online and has enough space by using "volume show" command.
3. Run the "lun import pause" command followed by a "lun import resume" command. In certain cases, the import might need to be restarted ( Use "lun import stop" command followed by "lun import start" command). The ultimate recovery is to delete this relationship (with -force option) and create it again.

### Syslog Message

Import of foreign LUN %s to LUN %s failed - Reason:%s

### Parameters

**DiskSerialNumber** (STRING): Foreign disk serial number.

**LUNSerialNumber** (STRING): Serial number of the LUN.

**FailureReason** (STRING): Internal error reason.

## fli.lun.import.inconsistent

### Severity

ERROR

### Description

This message occurs when the import inconsistent flag cannot be unset. The import inconsistent flag is a system attribute for the NetApp LUN that is set during import progress, to indicate inconsistent data in the destination LUN. This is to discourage restore operations on LUN Snapshot(tm) copies created during import. A best effort is done to unset this flag after the import is complete, and also during import stop and delete operations. It is safe to perform restore operations if data in the destination LUN is known to be successfully verified.

### Corrective Action

(None).

### Syslog Message

Import inconsistent flag could not be unset for LUN %s - Reason:%s

### Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**FailureReason** (STRING): Internal error reason.

## fli.lun.verify.complete

**Severity**

INFORMATIONAL

**Description**

This message occurs when a foreign LUN import verify is complete.

**Corrective Action**

(None).

**Syslog Message**

Import verify of foreign LUN %s of size %lld bytes from array model %s belonging to vendor %s with NetApp LUN %s is successfully completed.

**Parameters**

**DiskSerialNumber** (STRING): Foreign disk serial number.

**Size** (LONGINT): Size of the foreign LUN.

**ForeignArrayModelID** (STRING): Foreign array model ID.

**ForeignArrayVendor** (STRING): Foreign array name.

**LUNSerialNumber** (STRING): Serial number of the LUN.

**fli.lun.verify.failed****Severity**

ERROR

**Description**

This message occurs when a foreign LUN import verify fails. 1. Data could not be read from the source(foreign LUN). 2. Data could not be written to the destination(NetApp LUN). 3. There was a data mismatch during the verify operation. 4. There was an internal error.

**Corrective Action**

1. For a data mismatch, start the import operation again.
2. Verify the connectivity to the foreign disk.
3. Verify that destination volume is in a healthy state. In certain cases, the verify operation might need to be restarted (Use "lun import verify stop" command followed by "lun import verify start" command). The ultimate recovery is to delete this relationship (with -force option) and create it again.

**Syslog Message**

Import verify of foreign LUN %s and NetApp LUN %s failed - Reason:%s

**Parameters**

**DiskSerialNumber** (STRING): Foreign disk serial number.

**LUNSerialNumber** (STRING): Serial number of the LUN.

**FailureReason** (STRING): Internal error reason.

# fm events

## fm.cpuutilduringarl events

### fm.cpuUtilDuringARL

#### Deprecated

Deprecated as of version 9.0.

#### Severity

NOTICE

#### Description

This message occurs at the start of an aggregate relocation (ARL) and at the completion of an ARL. It records the maximum, minimum, and average CPU and disk utilization on the node executing the ARL.

#### Corrective Action

(None).

#### Syslog Message

CPU and disk utilization during the %d seconds %s: cpu\_util\_high: %lld; cpu\_util\_low: %lld; cpu\_util\_avg: %lld; disk\_util\_high: %lld; disk\_util\_low: %lld; disk\_util\_avg: %lld

#### Parameters

**window\_sz** (INT): Duration, in seconds, over which CPU and disk utilization are tracked.

**when** (STRING): Event during which CPU and disk utilization are tracked.

**cpu\_util\_high** (LONGINT): Maximum CPU utilization.

**cpu\_util\_low** (LONGINT): Minimum CPU utilization.

**cpu\_util\_avg** (LONGINT): Average CPU utilization.

**disk\_util\_high** (LONGINT): Maximum disk utilization.

**disk\_util\_low** (LONGINT): Minimum disk utilization.

**disk\_util\_avg** (LONGINT): Average disk utilization.

## fm.cpuutilduringtransition events

### fm.cpuUtilDuringTransition

#### Severity

NOTICE

#### Description

This message occurs during various failover and aggregate mobility events, including start of a takeover, end of a successful takeover, start of a CFO giveback, after completion of SFO giveback and completion of an aggregate relocation. This message records the maximum, minimum, and average CPU and disk utilization on the node executing the failover or aggregate mobility operation.

#### Corrective Action

(None).

## Syslog Message

CPU and disk utilization during the %d seconds %s: cpu\_util\_high: %lld; cpu\_util\_low: %lld; cpu\_util\_avg: %lld; disk\_util\_high: %lld; disk\_util\_low: %lld; disk\_util\_avg: %lld

## Parameters

**window\_sz** (INT): Duration, in seconds, over which CPU and disk utilization are tracked.

**when** (STRING): Event during which CPU and disk utilization are tracked.

**cpu\_util\_high** (LONGINT): Maximum CPU utilization.

**cpu\_util\_low** (LONGINT): Minimum CPU utilization.

**cpu\_util\_avg** (LONGINT): Average CPU utilization.

**disk\_util\_high** (LONGINT): Maximum disk utilization.

**disk\_util\_low** (LONGINT): Minimum disk utilization.

**disk\_util\_avg** (LONGINT): Average disk utilization.

# fm.overrideMbxUncertain events

## fm.overrideMbxUncertain

### Severity

NOTICE

### Description

This message occurs when the user has chosen to override a mailbox uncertain state using the CTRL\_O option. The node could have gotten into the mailbox uncertain state if, while booting up, it finds that one or more disks are missing and it does not have the requisite number of mailbox disks to boot up. If the override was successful, at least half of the online disks were also present in the mailbox hints bootarg and the node will boot up. If the override was unsuccessful, the node continues to remain in an uncertain mailbox state and will not boot up completely.

### Corrective Action

(None).

### Syslog Message

HA mailbox uncertain state overridden.

### Parameters

(None).



# fmmb events

## fmmb.add events

### fmmb.add.problem

#### Severity

ERROR

#### Description

This event is issued when the HA mailbox is attempting to add a mailbox disk and a problem occurred.

#### Corrective Action

(None).

#### Syslog Message

Error attempting to add disk %s as a %s HA mailbox disk.

#### Parameters

**disk** (STRING): The name of the disk. If the disk name is not available, a ?? will be used.

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

## fmmb.blobnotfound events

### fmmb.BlobNotFound

#### Severity

ALERT

#### Description

This message occurs when a mailbox 'blob' is not read from the on-disk mailbox data structure. A 'blob' is a piece of data which has been stored for a non-High Availability (HA) Data ONTAP® subsystem in the mailbox.

#### Corrective Action

If the blob id is "3" and this is a 2-node cluster, use the 'cluster ha modify -configure' command to unconfigure ha and then reconfigure it.

#### Syslog Message

(None).

#### Parameters

**blob\_id** (INT): 'blob' ID.

**owner** (STRING): Owner of the blob, either PARTNER or LOCAL.

**reason** (STRING): Reason for the problem.

## fmmb.blobnotwritten events

## fmmb.BlobNotWritten

### Deprecated

Deprecated as of version 9.2.

### Severity

ALERT

### Description

This message occurs when a mailbox 'blob' write operation to the on-disk mailbox data structure fails. A 'blob' is a piece of data that has been stored for a non-high availability (HA) Data ONTAP® subsystem in the HA mailbox.

### Corrective Action

If the blob ID is "3" and this is a two-node cluster, use the 'cluster ha modify -configure' command to unconfigure HA and then reconfigure it.

### Syslog Message

(None).

### Parameters

**blob\_id** (INT): 'blob' ID.

**owner** (STRING): Owner of the blob, either PARTNER or LOCAL.

**reason** (STRING): Reason for the problem.

## fmmb.blobwritten events

### fmmb.BlobWritten

### Severity

INFORMATIONAL

### Description

This message occurs when a mailbox 'blob' write to the on-disk mailbox data structure succeeds. A 'blob' is a piece of data that has been stored for a non-high availability (HA) Data ONTAP® subsystem in the HA mailbox.

### Corrective Action

(None).

### Syslog Message

Blob %d: bytes written %d: data: %llx.

### Parameters

**blob\_id** (INT): 'blob' ID.

**blob\_data\_len** (INT): Total number of bytes of data written.

**data** (LONGINTEX): First eight bytes of the data written.

# fmmb.current events

## fmmb.current.lock.disk

### Severity

INFORMATIONAL

### Description

This event is issued when HA mailbox disks are discovered or allocated.

### Corrective Action

(None).

### Syslog Message

Disk %s is a %s HA mailbox disk.

### Parameters

**disk** (STRING): The name of the disk. If the disk name is not available, the name ?? will be used.

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

# fmmb.destroy events

## fmmb.destroy.mbx

### Severity

INFORMATIONAL

### Description

This event is issued when the high-availability (HA) Mailbox contents are cleared, typically by the 'mailbox destroy' command.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk\_name** (STRING): Disk name on which the mailbox was cleared.

## fmmb.destroy.mbx.fail

### Severity

NOTICE

### Description

This event is issued when an attempt is made to destroy the mailbox contents on the disk and this fails. Data ONTAP to a 7.0 or later version.

**Corrective Action**

(None).

**Syslog Message**

Failed to clear mailbox on disk:%s, error:%s.

**Parameters**

**disk\_name** (STRING): disk name on which we are trying to clear the mailbox.

**error** (STRING): The reason why mailbox clear failed.

## fmmb.disk events

### fmmb.disk.IDchange

**Severity**

INFORMATIONAL

**Description**

This event is issued when the unique identifier of one of the disks used by the HA mailbox has changed.

**Corrective Action**

(None).

**Syslog Message**

Disk %s has changed its unique identifier.

**Parameters**

**disk** (STRING): The name of the disk whose identifier has changed.

### fmmb.disk.notAccsble

**Severity**

NOTICE

**Description**

This message occurs when mailbox disks of the specified host are not accessible.

**Corrective Action**

(None).

**Syslog Message**

All %s mailbox disks are inaccessible.

**Parameters**

**host** (STRING): Host string, which is either local or partner.

## fmmb.exthdr events

## fmmb.exthdr.null

### Severity

ERROR

### Description

This event is issued when the HA Mailbox extension header is invalid. The extension header is an on-disk structure which describes the opaque client blocks and major reset structures.

### Corrective Action

(None).

### Syslog Message

HA Mailbox extension header is null.

### Parameters

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

## fmmb.inststat events

### fmmb.instStat.change

### Severity

INFORMATIONAL

### Description

This message occurs whenever the mailbox instance state changes.

### Corrective Action

(None).

### Syslog Message

%s on %s side.

### Parameters

**buff** (STRING): Instance state string.

**host** (STRING): Host string, which is either local or partner.

## fmmb.lock events

### fmmb.lock.disk.remove

### Severity

INFORMATIONAL

### Description

This event is issued when a disk is removed from the HA mailbox disk lock set.

**Corrective Action**

(None).

**Syslog Message**

Disk %s removed from %s mailbox set.

**Parameters**

**disk** (STRING): The name of the disk. If the disk name is not available, a ?.? will be used.

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

## fmmb.major events

### fmmb.major.reset.invalid

**Severity**

ALERT

**Description**

This event is issued when the HA Mailbox major reset data should exist and it doesn't. The major reset data in the mailbox is created and used when there is a major change of information kept in the HA mailbox records.

**Corrective Action**

(None).

**Syslog Message**

HA %s Mailbox major reset header information is invalid.

**Parameters**

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

## fmmb.mediator events

### fmmb.mediator.lckdiskSuggest

**Severity**

INFORMATIONAL

**Description**

This message occurs when a mailbox disk from the mediator site is selected to be included in the mailbox disk list.

**Corrective Action**

(None).

**Syslog Message**

(None).

## Parameters

**side** (INT): The mailbox instance side. The value of 0 for local mailbox and the value of 1 for partner mailbox.

**uuid** (STRING): Disk unique ID of the mailbox disk on mediator site.

## fmmb.mediator.noLckdisk

### Severity

INFORMATIONAL

### Description

This message occurs when there is no mediator mailbox detected.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**side** (INT): The mailbox instance side. The value of 0 for local mailbox and the value of 1 for partner mailbox.

## fmmb.mediator.nullLckdiskSuggest

### Severity

INFORMATIONAL

### Description

This message occurs when there is null mediator mailbox detected.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**side** (INT): The mailbox instance side. The value of 0 for local mailbox and the value of 1 for partner mailbox.

## fmmb.mmr events

### fmmb.MMR.payload.bad

### Severity

ALERT

### Description

This event is issued when the outbound FM payload area is found to be corrupted and the master mailbox record is missing.

**Corrective Action**

(None).

**Syslog Message**

%s outbound FM payload area is invalid and the master mailbox record is missing.

**Parameters**

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

## fmmb.newdisksuggestion events

**fmmb.NewDiskSuggestion****Severity**

INFORMATIONAL

**Description**

This message occurs when RAID sends to the high-availability (HA) mailbox code a list of disks to use for HA mailbox disks.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**count** (INT): Number of suggested disks.

**new** (STRING): List of suggested disks.

**old** (STRING): Current list of HA mailbox disks.

## fmmb.oldformatcleared events

**fmmb.oldFormatCleared****Severity**

INFORMATIONAL

**Description**

This event is issued when the older format HA mailboxes (special area of disk used for bookkeeping) are cleared on an upgrade from a pre 7.0 version of Data ONTAP to a 7.0 or later version.

**Corrective Action**

(None).

**Syslog Message**

Clearing content of %s old format HA mailboxes (special area of disk used for bookkeeping).



## Parameters

**whichone** (STRING): Either local or partner to indicate which set of HA mailboxes are being cleared.

# fmmb.opaque events

## fmmb.opaque.hdr.invalid

### Severity

ALERT

### Description

This event is issued when the HA Mailbox opaque header is invalid. The opaque area in the HA Mailbox is used by HA mailbox clients, like RAID, to store persistent information.

### Corrective Action

(None).

### Syslog Message

HA %s Mailbox opaque header is invalid.

## Parameters

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

# fmmb.payload events

## fmmb.payload.outbound.bad

### Severity

ALERT

### Description

This event is issued when the outbound FM payload area is found to be corrupted. The invalid fm state information is logged. As a result of this condition the payload is zeroed and HA mailbox re-initialization is done.

### Corrective Action

(None).

### Syslog Message

%s outbound FM payload area is invalid with fmState of %d.

## Parameters

**host** (STRING): The node which owns the HA mailbox disk. The node can be either local or partner.

**fmState** (INT): The invalid FM state information.

# fmmb.revert events

## **fmmb.revert.opaque**

### **Severity**

ERROR

### **Description**

This event is issued when the HA Mailbox opaque receives a write request for the opaque area and a revert\_to command is in progress. The write request is rejected. The opaque area in the HA Mailbox is used by non HA mailbox clients, like RAID, to store persistent information.

### **Corrective Action**

(None).

### **Syslog Message**

HA Mailbox RAID write request rejected during revert.

### **Parameters**

(None).

# fp events

## fp.est events

### fp.est.scan.catalog.failed

#### Severity

ERROR

#### Description

This message occurs when the volume footprint estimator scan fails to update the catalog file. This can cause footprint data reduction to be unreported post volume remount until the next scan updates the file.

#### Corrective Action

Check if there is atleast 4K worth of space available for the catalog metafile in the volume. Use "df <vol-name>" command to see available space details.

#### Syslog Message

Volume footprint estimator scan catalog update fails on "%s" - %s.

#### Parameters

**vol** (STRING): Volume name.

**fail\_message** (STRING): Update failure reason.

### fp.est.scan.catalog.updated

#### Severity

NOTICE

#### Description

This message occurs when related information of the volume footprint estimator scan is updated on the catalog file.

#### Corrective Action

(None).

#### Syslog Message

Volume footprint estimator scan catalog updated on '%s' (SE ratio local: %s, SE ratio migrate: %s, Incompressible: %d%%, Hot data: %llu%%, Mode: %d, SE ratio cold migrate: %s).

#### Parameters

**vol** (STRING): Volume name.

**se\_ratio\_local** (STRING): Storage efficiency ratio for local reporting.

**se\_ratio\_migrate** (STRING): Storage efficiency ratio for migrate reporting.

**incomp\_data\_pct** (INT): Incompressible data percentage.

**hot\_data\_pct** (LONGINT): Hot data percentage.

**scan\_mode** (INT): Scan mode.

**se\_ratio\_cold\_migrate** (STRING): Storage efficiency ratio for cold migrate reporting.

## **fp.est.scan.start.failed**

### **Severity**

ERROR

### **Description**

This message occurs when the volume footprint estimator scan fails to start. This might cause footprint data reduction reporting to be inaccurate. But the scan is expected to be triggered again in the next cycle which is likely to be the next day.

### **Corrective Action**

Scan will be triggered again in the next cycle which is likely to be the next day.

### **Syslog Message**

Volume footprint estimator scan fails to start on "%s" - %s.

### **Parameters**

**vol** (STRING): Volume name.

**fail\_message** (STRING): Scan failure reason.

# fpolicy events

## fpolicy.eagain events

### fpolicy.eagain.disconnect

#### Severity

ALERT

#### Description

This message occurs when the TCP connection to the FPolicy server is broken four times in a two minute interval due to EAGAIN error. ONTAP will not establish a new connection and FPolicy notifications will not be sent to the server.

#### Corrective Action

Persistent EAGAIN errors indicate slow network or slow/unresponsive FPolicy server. Check the connection between the FPolicy server and the node or contact the FPolicy application's vendor support to address the slow/unresponsive server.

#### Syslog Message

Connection to the FPolicy server "%s" of policy "%s" is broken for Vserver %s ( reason: "%s" ).

#### Parameters

**fpserver** (STRING): IP address of the FPolicy server.

**policy** (STRING): Name of the policy.

**vserver** (STRING): Vserver Name.

**reason** (STRING): Reason for disconnection.

### fpolicy.eagain.on.write

#### Severity

NOTICE

#### Description

This message occurs when there is no space in the send buffer for sending a message to the FPolicy server. The FPolicy server is apparently slow in reading notifications.

#### Corrective Action

Occasional EAGAIN (resource temporarily unavailable) errors can occur for transient network issues. If this message occurs frequently, check the connection between the FPolicy server and the node. If the errors persist, contact the FPolicy application's vendor support.

#### Syslog Message

Write returned EAGAIN while sending notification to the FPolicy server "%s" for Vserver %s with UUID %s.

#### Parameters

**fpserver** (STRING): IP address of the FPolicy server.

**vserver** (STRING): Vserver Name.

**Vserver\_UUID** (STRING): Vserver UUID.

# fpolicy.fgconversion events

## fpolicy.fgconversion.warning

### Severity

NOTICE

### Description

This message occurs when a FlexVol® unmount notification is not sent to the FPolicy(tm) server during a FlexGroup conversion.

### Corrective Action

(None).

### Syslog Message

Failed to send FPolicy unmount notification for FlexVol %s of Vserver %s. FlexGroup MSID: %llu.

### Parameters

**vserverName** (STRING): Name of the Vserver.

**volumeName** (STRING): Name of the volume being converted to a FlexGroup.

**flexgroupMsid** (LONGINTEX): MSID of the new FlexGroup.

# fpolicy.memory events

## fpolicy.memory.pressure

### Severity

ALERT

### Description

This message occurs when either the FPolicy Service Manager (FSM) receives too many requests to process immediately, or when the number of requests is accumulating due to server connectivity issues. FSM processing time might degrade as a result.

### Corrective Action

Check the connection between the node and the FPolicy server by using the "vserver fpolicy show-engine" command. Verify that FPolicy is only monitoring file operations that are recommended by the FPolicy server vendor. Verify that FPolicy is using filters recommended by the server vendor for common file operations, for example, close, read, write and open.

### Syslog Message

FSM is experiencing memory pressure.

### Parameters

(None).

# fpolicy.outage events

## **fpolicy.outage.file.invalid**

### **Severity**

ERROR

### **Description**

This message occurs when file creation fails under the configured resiliency directory path.

### **Corrective Action**

Check the configured resiliency directory path. To find the configured resiliency directory path by using the "vserver fpolicy policy external-engine show" command.

### **Syslog Message**

Failed to create file under configured resiliency directory path of FPolicy external engine "%s" for policy "%s". Reason: %s.

### **Parameters**

**engine** (STRING): FPolicy external engine.

**policy** (STRING): FPolicy policy.

**reason** (STRING): Reason for the failure.

## **fpolicy.outage.msg.lost**

### **Severity**

ERROR

### **Description**

This message occurs when the send operation fails for the network outage notification.

### **Corrective Action**

Find the file in the configured resiliency directory path.

### **Syslog Message**

Failed to send outage notification message to any FPolicy server for policy "%s" and path "%s".

### **Parameters**

**policy** (STRING): FPolicy policy.

**path** (STRING): File path.

## **fpolicy.outage.msg.success**

### **Severity**

NOTICE

### **Description**

This message occurs when a network outage notification is successfully sent.

### **Corrective Action**

(None).

## Syslog Message

Outage notification message sent to at least one FPolicy server for policy "%s" and file path "%s".

## Parameters

**policy** (STRING): FPolicy policy.

**path** (STRING): File path.

## fpolicy.outage.over

### Severity

NOTICE

### Description

This message occurs when the FPolicy external engine, acting asynchronously and with resiliency support enabled, is reconnected after a network outage.

### Corrective Action

(None).

## Syslog Message

The network outage (duration: "%llu" sec) is over. At least one FPolicy server is now connected successfully for FPolicy external engine "%s" of policy "%s".

## Parameters

**Duration** (LONGINT): Network outage duration in seconds.

**engine** (STRING): Asynchronous FPolicy external engine.

**policy** (STRING): FPolicy policy.

## fpolicy.outage.start

### Severity

NOTICE

### Description

This message occurs when the FPolicy external engine, acting asynchronously and with resiliency support enabled, is disconnected due to a network outage.

### Corrective Action

(None).

## Syslog Message

Network outage for FPolicy external engine "%s" of policy "%s".

## Parameters

**engine** (STRING): Asynchronous FPolicy external engine.

**policy** (STRING): FPolicy policy.

## fpolicy.outage.timeout



**Severity**

NOTICE

**Description**

This message occurs when the maximum notification retention duration for the FPolicy external engine expires, and the engine is configured to operate asynchronously with resiliency support enabled. The storage controller stores no more notifications.

**Corrective Action**

(None).

**Syslog Message**

Maximum notification retention time has expired for FPolicy external engine "%s" of policy "%s".

**Parameters**

**engine** (STRING): Asynchronous FPolicy external engine.

**policy** (STRING): FPolicy policy.

**fpolicy.outage.write****Severity**

NOTICE

**Description**

This message occurs when FPolicy cannot store file notifications using the configured resiliency directory path.

**Corrective Action**

(None).

**Syslog Message**

Failed to store notifications in configured resiliency directory path of FPolicy external engine "%s" for policy "%s". Reason: %s.

**Parameters**

**engine** (STRING): FPolicy external engine.

**policy** (STRING): FPolicy policy.

**reason** (STRING): Reason for the failure.

**fpolicy.outage.writes.lost****Severity**

ERROR

**Description**

This message occurs when FPolicy fails multiple consecutive times to store file notifications using the configured resiliency directory path.

**Corrective Action**

Check the configured resiliency directory path. To find the configured resiliency directory path by using the "vserver fpolicy policy external-engine show" command.

## Syslog Message

Unable to store notifications in configured resiliency directory path of FPolicy external engine "%s" for policy "%s".

## Parameters

**engine** (STRING): FPolicy external engine.  
**policy** (STRING): FPolicy policy.

# fpolicy.pspath events

## fpolicy.pspath.failed

### Severity

ERROR

### Description

This message occurs when the creation of the FPolicy persistent store active or backlog directories fails. This disrupts the persistent store workflow in FPolicy.

### Corrective Action

Check the space availability and permission on volume and recreate the PS config.

## Syslog Message

Failed to create FPolicy persistent store path "%s" in SVM "%s".

## Parameters

**path** (STRING): The directory path name that failed while it's creation.  
**vserverName** (STRING): Name of the Vserver.

# fpolicy.server events

## fpolicy.server.bkPressApply

### Severity

INFORMATIONAL

### Description

This message occurs when the FPolicy server has applied backpressure which means that the server is busy and does not want to receive requests until it removes backpressure.

### Corrective Action

(None).

## Syslog Message

FPolicy server "%s" has applied backpressure.

## Parameters

**fpserver** (STRING): IP Address of the FPolicy server.

## **fpolicy.server.bkPressRemove**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the FPolicy server removes backpressure, which it has previously applied. The FPolicy server removes backpressure when it is ready to receive requests.

### **Corrective Action**

(None).

### **Syslog Message**

FPolicy server "%s" has removed backpressure.

### **Parameters**

**fpserver** (STRING): IP Address of the FPolicy server.

## **fpolicy.server.connect**

### **Severity**

NOTICE

### **Description**

This message occurs when a control channel is established between the node and the FPolicy server.

### **Corrective Action**

(None).

### **Syslog Message**

Control channel with the FPolicy server "%s" of policy "%s" is established for Vserver %s.

### **Parameters**

**fpserver** (STRING): IP Address of the FPolicy server.

**policy** (STRING): Name of the policy.

**vserver** (STRING): Vserver Name.

## **fpolicy.server.connectError**

### **Severity**

ERROR

### **Description**

This message occurs when the node fails to establish a control channel with the FPolicy server. The control channel is used for communication between the node and the FPolicy server. The error could be because of a failure in establishing TCP connection, SSL failure, FPolicy handshake failure, or some internal error.

### **Corrective Action**

Depending on the failure reason, check the following: On the node, check whether the correct IP and port entries are specified in the policy configuration. Also, the FPolicy server should be running and listening on the port specified in the policy configuration. Run the command "network interface show" to determine that at least one data logical interface (LIF) is configured to the Vserver on the node. Run the command

"routing-groups route show" to verify that the network route exists to the FPolicy server. Run the command "network ping" to verify that the FPolicy server is reachable from the node. For any other reason, see the Cluster-Mode documentation.

### Syslog Message

Node failed to establish a connection with the FPolicy server "%s" of policy "%s" for Vserver %s (reason: "%s").

### Parameters

**fpserver** (STRING): IP address of the FPolicy server.  
**policy** (STRING): Name of the policy.  
**vserver** (STRING): Vserver Name.  
**reason** (STRING): Reason for failure.

## fpolicy.server.disconnect

### Severity

ERROR

### Description

This message occurs when the TCP connection to the FPolicy(tm) server is broken.

### Corrective Action

Reestablish the connection by using "vserver fpolicy engine-connect" command.

### Syslog Message

Connection to the FPolicy server "%s" of policy "%s" is broken for Vserver %s ( reason: "%s" ).

### Parameters

**fpserver** (STRING): IP address of the FPolicy server.  
**policy** (STRING): Name of the policy.  
**vserver** (STRING): Vserver Name.  
**reason** (STRING): Reason for disconnection.

## fpolicy.server.maxPendingReq

### Severity

INFORMATIONAL

### Description

This message occurs when the maximum outstanding requests for a FPolicy server from a node has reached. The value is configurable through CLI.

### Corrective Action

If you want to increase the maximum pending request modify the "-max-server-reqs" option using command "fpolicy policy external-engine modify" from CLI.

### Syslog Message

Maximum outstanding requests for the FPolicy server "%s" has reached.

## Parameters

**fpserver** (STRING): IP Address of the FPolicy server.

# fsm events

## fsm.low events

### fsm.low.memory

#### Severity

ERROR

#### Description

This message occurs when Fpolicy Service Manager (FSM) fails to send outage, alert, handshake, cancel or volume operation notification to the FPolicy server due to memory allocation error.

#### Corrective Action

Check the connection between ONTAP and FPolicy server. Slow connection may result in accumulation of notifications in FSM. Make sure that the FPolicy is monitoring file operation which are recommended by FPolicy server vendor. Also use vendor recommended filter for operation like close, read, write and open.

#### Syslog Message

Failed to allocate memory for %s notification.

#### Parameters

**reqType** (STRING): Type of the message.

# gb events

## gb.cfo events

### gb.cfo.abort.raid.fm

#### Severity

ERROR

#### Description

This message occurs when a high-availability (HA) controller failover (CFO) giveback is canceled due to RAID activity on a partner aggregate.

#### Corrective Action

Reissue the "storage failover giveback" command to initiate giveback to the partner after the pending or in-progress operation on the partner volume or aggregate is complete. Use the "storage aggregate status -r" command to view the state of the aggregate. If a plex is offline, use the "storage aggregate plex online" command to bring the plex online. If unmirrored aggregates are not supported on this configuration, use the "storage aggregate mirror" command to mirror this aggregate.

#### Syslog Message

%s %s%s is %s; canceling giveback.

#### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Current owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**reason** (STRING): Activity that is preventing CFO giveback.

### gb.cfo.veto.asup.generalNopart

#### Severity

NOTICE

#### Description

This message occurs when AutoSupport vetoes a controller failover (CFO) giveback request because an AutoSupport message for the down partner node was not collected prior to the giveback request.

#### Corrective Action

Typically, an AutoSupport message for the partner node is collected shortly after takeover is complete. If the CFO giveback was attempted and vetoed within 20 minutes of the takeover, retry the operation later. If time does not permit waiting, use the 'cf giveback -f' command to force a giveback.

#### Syslog Message

AutoSupport vetoed giveback because an AutoSupport message for the down partner node was not yet collected.

#### Parameters

(None).

## **gb.cfo.veto.coredump.pardump**

### **Severity**

NOTICE

### **Description**

This message occurs when the partner node is still dumping core, and cannot yet resume serving data. The giveback attempt has been canceled.

### **Corrective Action**

Either let the dumpcore operation finish, or use the "-override-vetoes true" parameter to force a giveback.

### **Syslog Message**

Giveback canceled because partner is still dumping core

### **Parameters**

(None).

## **gb.cfo.veto.coredump.parsave**

### **Severity**

NOTICE

### **Description**

This message occurs when savecore is currently saving the partner coredump. This work will be lost if a giveback is done at this time, so the giveback has been canceled.

### **Corrective Action**

Either let the savecore operation finish, or use the "-override-vetoes true" parameter to force a giveback.

### **Syslog Message**

Giveback canceled because savecore is running. %d%% - %d of %d blocks saved

### **Parameters**

**percent** (INT): Percentage of core that has been saved  
**completed** (INT): Number of blocks that have been saved  
**total** (INT): Total number of blocks in the core

## **gb.cfo.veto.coredump.savinit**

### **Severity**

NOTICE

### **Description**

This message occurs when savecore is getting ready to start saving the partner coredump. The giveback attempt has been canceled.

### **Corrective Action**

Either let the savecore operation finish, or use the "-override-vetoes true" parameter to force a giveback.



## Syslog Message

Giveback canceled because savecore is initializing

## Parameters

(None).

## gb.cfo.veto.lmgr.nonCA.broke

### Severity

NOTICE

### Description

This message occurs when a giveback is forced that causes non-continuously available (non-CA) locks on the volume to be dropped. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in controller failover (CFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

### Corrective Action

(None).

## Syslog Message

Dropped non-CA locks on volume %s%s%s%s on CFO aggregate %s due to forced giveback.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

**aggrname** (STRING): Aggregate name.

## gb.cfo.veto.lmgr.nonCA.locks

### Severity

ERROR

### Description

This message occurs when a giveback has started but cannot proceed because non-continuously available (non-CA) locks are present on the volume. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in controller failover (CFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

### Corrective Action

Based on how resilient they are to failures, applications must either gracefully close sessions over which non-CA locks are established or accept lock state disruption. To determine the open files that have these sessions established, run the 'vserver cifs session file show -hosting-aggregate "aggregate list" -continuously-available No' command. "aggregate list" is a list of aggregates that are sent home as a result of the giveback operation. If lock state disruption for all existing non-CA locks is acceptable, retry the giveback operation by using the '-override-vetoes true' option.

## Syslog Message

Could not complete giveback because of non-CA locks on volume %s%s%s%s CFO aggregate %s.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

**aggrname** (STRING): Aggregate name.

## gb.cfo.veto.nvsave.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform a giveback due to WAFL® NVSave file replay activity, and must abort the giveback.

### Corrective Action

Retry the giveback operation after completion of NVSave files replay. Replay completion is indicated by the EMS message 'wafl.nvsave.replaying.pass' or 'wafl.nvsave.replaying.fail'.

## Syslog Message

%s %s%s%s%s has WAFL NVSave file replay activity underway; canceling giveback.

## Parameters

**type** (STRING): Object type (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## gb.cfo.veto.repl

### Severity

NOTICE

### Description

This message occurs when a cluster failover(CFO) giveback is vetoed by a critical replication transfer (for example, volume move in the cutover phase).

### Corrective Action

Either abort the replication transfer that is preventing the giveback from being completed, retry the 'storage failover giveback' command with the '-override-vetoes true' option, or wait for the transfer to finish.

## Syslog Message

Could not complete giveback because a replication transfer with UUID %s is in progress.

## Parameters

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed the giveback.

## gb.cfo.veto.sanown.ReassignInProgress

### Severity

ERROR

### Description

This message occurs when a controller disallows a giveback attempt because the partner's disk ownership is being reassigned. Typically, ownership reassignment is done after a controller replacement.

### Corrective Action

Reissue the giveback request after disk ownership assignment is completed.

### Syslog Message

Giveback was vetoed because the disk ownership reassignment in progress on the partner cannot be interrupted.

### Parameters

(None).

## gb.cfo.veto.snap.fail

### Severity

NOTICE

### Description

This message occurs when the system cannot perform a giveback due to Snapshot(tm) copy activity, and must abort the giveback.

### Corrective Action

(None).

### Syslog Message

%s %s%s%s%s has Snapshot copy activity underway; canceling giveback.

### Parameters

**type** (STRING): Object type (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the the volume name itself is insufficient.

## gb.cfo.veto.snap.warning

### Severity

NOTICE

### Description

This message occurs when the system detects Snapshot(tm) copy activity during giveback.

## Corrective Action

(None).

## Syslog Message

%s %s%s%s%s has Snapshot copy activity underway.

## Parameters

**type** (STRING): Object type (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the the volume name itself is insufficient.

## gb.cfo.veto.storageDiskBrkn

### Severity

ERROR

### Description

This message occurs when a giveback is aborted because a failed disk is present.

### Corrective Action

Remove all failed disks and retry the giveback.

### Syslog Message

Failed disk %s should be removed before the giveback command is invoked.

### Parameters

**diskName** (STRING): Name of the device.

## gb.cfo.veto.storageDiskByp

### Severity

ERROR

### Description

This message occurs when a giveback is aborted because one or more bypassed disks are present.

### Corrective Action

Remove all bypassed disks and retry the giveback.

### Syslog Message

All bypassed disks should be removed before a giveback command is entered.

### Parameters

(None).

## gb.cfo.veto.storageDiskSanit

**Severity**

ERROR

**Description**

This message occurs when a giveback is aborted because disk sanitization is in progress.

**Corrective Action**

Retry the giveback after sanitization is completed.

**Syslog Message**

Disk sanitization is in the process of formatting a disk. This operation should not be interrupted by a giveback.

**Parameters**

(None).

**gb.cfo.veto.vservermigrate****Severity**

ERROR

**Description**

This message occurs when the vserver migrate cutover operation vetoes cfo aggregate giveback because vserver migration cannot be aborted during that time.

**Corrective Action**

Reissue the 'storage failover giveback' command to initiate giveback to the partner when the vserver migrate cutover operation is completed.

**Syslog Message**

Vserver migration in progress prevents giveback.

**Parameters**

(None).

**gb.netra events****gb.netra.ca.check.failed****Severity**

ERROR

**Description**

This message occurs during the migration of an aggregate as a part of storage failover (SFO) giveback, when the destination node cannot reach the object stores.

**Corrective Action**

Verify that your intercluster LIF is online and functional by using the 'network interface show' command. Also, check network connectivity to the object store server by using the 'ping' command over the destination node intercluster LIF. Additionally, verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the 'aggregate object-store config show'

command. It is possible to override this error by specifying false for the "require-partner-waiting" parameter of the giveback command. For more information or assistance, contact NetApp technical support.

### Syslog Message

Giveback of aggregate '%s' (uuid: %s) failed due to %s preventing object store access on the destination node.

### Parameters

**vol** (STRING): Name of the aggregate.

**aggr\_uuid** (STRING): UUID of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate giveback operation.

## gb.sfo events

### gb.sfo.abort.raid.fm

#### Severity

ERROR

#### Description

This message occurs when a high-availability (HA) storage failover (SFO) giveback is canceled due to RAID activity on a partner aggregate.

#### Corrective Action

Reissue the "storage failover giveback" command to initiate giveback to the partner after the pending or in-progress operation on the partner volume or aggregate is complete. Use the "storage aggregate status -r" command to view the state of the aggregate. If a plex is offline, use the "storage aggregate plex online" command to bring the plex online. If unmirrored aggregates are not supported on this configuration, use the "storage aggregate mirror" command to mirror this aggregate.

#### Syslog Message

%s %s%s is %s; canceling giveback.

#### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Current owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing SFO giveback.

### gb.sfo.netra.ha.dskChkFailed

#### Severity

ERROR

#### Description

This message occurs when giveback of the partner's storage failover (SFO) aggregate fails because the high-availability (HA) partner cannot see all disks belonging to the aggregate.

#### Corrective Action

Use the "storage failover show -fields local-missing-disks, partner-missing-disks" command to rescan the disks and to determine the latest ownership. If the disk inventory mismatch is not resolved, check the

cabling for loose connections or a bad cable on the partner node.

### **Syslog Message**

Giveback of aggregate %s failed because the partner cannot see all disks belonging to the aggregate.

### **Parameters**

**aggr** (STRING): Name of the aggregate that was not given back.

## **gb.sfo.netra.raid.failed**

### **Severity**

ERROR

### **Description**

This message occurs during migration of an aggregate as part of storage failover (SFO) giveback, when one of the aggregate checks on the destination node fails, including duplicate aggregate name or duplicate Universally Unique Identifier (UUID).

### **Corrective Action**

1. If checks fail due to duplicate aggregate name or UUID, delete the duplicate aggregate created from previous "storage failover giveback" or "storage aggregate relocation" operations. Contact NetApp technical support for assistance with deletion of the unwanted duplicate aggregate.

### **Syslog Message**

Aggregate '%s' (UUID = %s, type = %s, home\_owner\_id = %llu, dr\_home\_owner\_id = %llu): RAID aggregate migration checks failed on destination node %s.

### **Parameters**

**vol** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): UUID of the aggregate.

**raid\_type** (STRING): RAID type of the volume.

**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.

**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's disaster recovery (DR) home owner.

**reason** (STRING): Activity that is preventing the aggregate relocation or giveback operation.

## **gb.sfo.netra.wafl.mcc.veto**

### **Severity**

ERROR

### **Description**

This message occurs during migration of an aggregate as part of storage failover (SFO) giveback, when one or more online left-behind disaster recovery(DR) aggregates are found on the destination node and the "node-object-limit" option is set to off.

### **Corrective Action**

Check whether there exists any online left-behind DR aggregates on the destination node from a previous switchover operation. If such an aggregate exists, then perform the corrective actions specified in the earlier EMS messages for the left-behind aggregate to return it to its original owner, and then retry the operation. If you cannot perform the corrective action, then use the "override-vetoes" option in the "giveback" command to force the giveback.

### **Syslog Message**

Aggregate '%s': One or more online left-behind DR aggregates were found on the destination node and the "node-object-limit" option is set to off.

### **Parameters**

**vol** (STRING): Name of the aggregate.

## **gb.sfo.precmt.lmgr.resynabrt**

### **Severity**

NOTICE

### **Description**

This message occurs during giveback, when automatic lock resynchronization with the node's high-availability (HA) partner for files on this aggregate was aborted.

### **Corrective Action**

(None).

### **Syslog Message**

Automatic lock resynchronization was aborted for aggregate %s on forced giveback.

### **Parameters**

**aggr** (STRING): Name of the aggregate for which resynchronization was aborted.

## **gb.sfo.precmt.lmgr.resyncing**

### **Severity**

ERROR

### **Description**

This message occurs during giveback, when automatic lock resynchronization with the node's high-availability (HA) partner is in progress for files on this aggregate.

### **Corrective Action**

Retry the giveback after verifying that lock synchronization is complete using the 'debug locks auto-resync-status -aggregate <aggregate>' diagnostic privilege command. If lock state disruption for all CA locks on the aggregate is acceptable, retry the giveback with the '-override-vetoes true' option.

### **Syslog Message**

Could not give back aggregate %s because automatic lock resynchronization was in progress.

### **Parameters**

**aggr** (STRING): Name of the aggregate with locks not yet in sync.

## **gb.sfo.precmt.repl**

### **Severity**

NOTICE



**Description**

This message occurs when all replication transfers are aborted because of storage failover (SFO) giveback.

**Corrective Action**

(None).

**Syslog Message**

Replication transfers aborted due to SFO giveback.

**Parameters**

(None).

**gb.sfo.precmt.vdom.lowMem****Severity**

ERROR

**Description**

This message occurs when giveback of the partner's storage failover (SFO) aggregate fails due to the unavailability of Data ONTAP® memory on the source node.

**Corrective Action**

Free some memory and retry the giveback.

**Syslog Message**

Giveback of aggregate %s failed due to unavailability of Data ONTAP® memory on the source node.

**Parameters**

**aggr** (STRING): Name of the aggregate that was not given back.

**gb.sfo.precmt.wafl.volconversion****Severity**

ERROR

**Description**

This message occurs when an active volume conversion is running on one or more volumes on the source aggregate. Pre-commit is aborted.

**Corrective Action**

When volume conversion is complete, re-run the 'storage failover giveback' command.

**Syslog Message**

SFO giveback is blocked as Volume Conversion is in progress on one of the volumes on aggregate %s.

**Parameters**

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

## **gb.sfo.veto.asup.generalNopart**

### **Severity**

NOTICE

### **Description**

This message occurs when AutoSupport vetoes a storage failover (SFO) giveback (sendhome) request because an AutoSupport message for the down partner node was not collected prior to the giveback request.

### **Corrective Action**

Typically, an AutoSupport message for the partner node is collected shortly after takeover is complete. If an SFO giveback was attempted and vetoed within 20 minutes of the takeover, retry the operation later. If time does not permit waiting, use the 'storage failover giveback -override-vetoes true' command to force a giveback.

### **Syslog Message**

AutoSupport vetoed giveback because an AutoSupport message for the down partner node was not yet collected.

### **Parameters**

(None).

## **gb.sfo.veto.dump**

### **Severity**

ERROR

### **Description**

This message occurs when dump/restore is aborted during storage failover/ aggregate relocation.

### **Corrective Action**

(None).

### **Syslog Message**

Dump/restore operations were aborted because %s is in progress.

### **Parameters**

**arl\_or\_sendhome** (STRING): Describes whether ARL or SFO

## **gb.sfo.veto.kmgr.keysmismissing**

### **Severity**

ERROR

### **Description**

This message occurs when a giveback of the partner's storage failover (SFO) aggregate fails due to the unavailability of volume encryption keys for the encrypted volumes of the aggregate on the partner node.

### **Corrective Action**

Wait a few minutes, and then try the giveback of the partner's SFO aggregate again. If the problem persists, run the "security key-manager external restore" command for an external key manager, the "security key-

manager external azure restore" command for an AKV configuration, the "security key-manager external aws restore" command for an AWSKMS configuration, the "security key-manager external gcp restore" command for a GCPKMS configuration, the "security key-manager external ikp restore" command for an IKPKMS configuration, or the "security key-manager onboard sync" command for the Onboard Key Manager. These commands retrieve the encryption keys for the encrypted volumes of the aggregate from the key manager. After running the appropriate command, then try the giveback of the partner's SFO aggregate again.

### Syslog Message

Giveback of aggregate "%s" failed due to the unavailability of the volume encryption keys for the encrypted volumes of the aggregate on partner node "%s".

### Parameters

**aggr** (STRING): Name of the aggregate that was not given back.  
**node** (STRING): Name of the node that is missing encryption keys.

## gb.sfo.veto.lmgr.nonCA.broke

### Severity

NOTICE

### Description

This message occurs when a giveback is forced that causes non-continuously available (non-CA) locks on the volume to be dropped. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

### Corrective Action

(None).

### Syslog Message

Dropped non-CA locks on volume %s%s%s%s on SFO aggregate %s due to forced giveback.

### Parameters

**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.  
**aggrname** (STRING): Aggregate name.

## gb.sfo.veto.lmgr.nonCA.locks

### Severity

ERROR

### Description

This message occurs when a giveback has started but cannot proceed because non-continuously available (non-CA) locks are present on the volume. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

## Corrective Action

Based on how resilient they are to failures, applications must either gracefully close sessions over which non-CA locks are established or accept lock state disruption. To determine the open files that have these sessions established, run the 'vserver cifs session file show -hosting-aggregate "aggregate list" -continuously-available No' command. "aggregate list" is a list of aggregates that are sent home as a result of the giveback operation. If lock state disruption for all existing non-CA locks is acceptable, retry the giveback operation by using the '-override-vetoes true' option.

## Syslog Message

Could not complete giveback because of non-CA locks on volume %s%s%s%s SFO aggregate %s.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

**aggrname** (STRING): Aggregate name.

## gb.sfo.veto.lmgr.recons.left

### Severity

ERROR

### Description

This message occurs when a storage failover giveback has started, but cannot proceed because reconstruction of client file locks from mirrored lock information is not yet complete.

## Corrective Action

Retry the giveback after checking the status of reconstruction using the 'debug locks reconstruction show' diagnostic privilege command. If the locks reconstruction operation is not completed, restart the giveback and override the veto using the '-override-vetoes true' option.

## Syslog Message

Could not perform giveback because lock reconstruction on node %s is in progress.

## Parameters

**node** (STRING): Name of the node reconstructing locks.

## gb.sfo.veto.lmgr.syncing

### Severity

ERROR

### Description

This message occurs when a storage failover giveback has started but cannot proceed because synchronization of client file locks ("locks sync") with the partner is not yet complete.

## Corrective Action

Retry the giveback after verifying that lock synchronization is completed, by using the 'debug locks lock-sync show' diagnostic privilege command.

## Syslog Message

Could not perform giveback because locks sync from node %s to its partner node %s is in progress.

## Parameters

**node** (STRING): Name of the node performing locks sync with its partner.

**partner** (STRING): Name of the partner node with which locks sync is being performed.

## gb.sfo.veto.repl

### Severity

NOTICE

### Description

This message occurs when a storage failover(SFO) giveback is vetoed by a critical replication transfer (for example, volume move in the cutover phase).

### Corrective Action

Either abort the replication transfer that is preventing the giveback from being completed, retry the 'storage failover giveback' command with the '-override-vetoes true' option, or wait for the transfer to finish.

## Syslog Message

Could not complete giveback because a replication transfer with UUID %s involving the aggregate %s is in progress.

## Parameters

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed the giveback.

**aggr** (STRING): Name of the aggregate.

## gb.sfo.veto.snaprestore

### Severity

ERROR

### Description

This message occurs when the system cannot perform a giveback because a volume snaprestore operation is active. The giveback is aborted.

### Corrective Action

Retry giveback after the snaprestore operation is finished.

## Syslog Message

%s %s%s%s%s is running snaprestore. Canceling giveback.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## **gb.sfo.veto.volmove**

### **Severity**

ERROR

### **Description**

This message occurs when the volume move operation vetoes a giveback because volume move can not be aborted during that time.

### **Corrective Action**

Reissue the 'storage failover giveback' command to initiate giveback to the partner when the volume move is completed. Or use the '-override-vetoes true' option to abort the volume move.

### **Syslog Message**

Volume move in progress for volume %s with DSID %llu prevents giveback.

### **Parameters**

**vol\_name** (STRING): Name of the source or destination volume.

**volume\_dsid** (LONGINT): Source or destination volume data set identifier (DSID).

## **gb.sfo.veto.vservermigrate**

### **Severity**

ERROR

### **Description**

This message occurs when the vserver migrate cutover operation vetoes data aggregate giveback because vserver migration cannot be aborted during that time.

### **Corrective Action**

Reissue the 'storage failover giveback' command to initiate giveback to the partner when the vserver migrate cutover operation is completed.

### **Syslog Message**

Vserver migration in progress prevents giveback.

### **Parameters**

(None).

## **gb.sfo.veto.wafl.volconversion**

### **Severity**

ERROR

### **Description**

This message occurs when an active volume conversion against one or more of the volumes on the source aggregate vetoed giveback.

### **Corrective Action**

When volume conversion is complete, re-run the 'storage failover giveback' command.

### **Syslog Message**

SFO giveback is vetoed as Volume Conversion is in progress on one of the volumes on aggregate %s.

### **Parameters**

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

# geom events

## geom.ontap events

### geom.ontap.orphan.removing

#### Severity

NOTICE

#### Description

This message occurs when the geometry layer (geom) detects a device to be detached from the system.

#### Corrective Action

(None).

#### Syslog Message

Removing unit %d type %d.

#### Parameters

**unit** (INT): Unit number.

**disk\_type** (INT): Disk type.



# gpo events

## gpo.cap events

### gpo.CAP.applied

**Severity**

INFORMATIONAL

**Description**

This message occurs when applying group policy object (GPO) central access policy (CAP) settings succeeds.

**Corrective Action**

(None).

**Syslog Message**

Successfully applied CAP settings to Vserver %s.

**Parameters**

**vserver** (STRING): Vserver name.

### gpo.CAP.notApplied

**Severity**

ERROR

**Description**

This message occurs when applying group policy object (GPO) central access policy (CAP) settings fails.

**Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

**Syslog Message**

Failed to apply CAP settings to Vserver %s failed, reason: %s.

**Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.frmwork events

### gpo.frmwork.RSoPApplied

**Severity**

NOTICE

**Description**

This event occurs after the resultant set of Group Policy (GPO) is calculated and applied to the Replicated Database (RDB) successfully.

**Corrective Action**

(None).

**Syslog Message**

GPO policy applied into RDB successfully.

**Parameters**

**vserver** (STRING): Vserver name.

**gpo.frmwork.secJobFailed****Severity**

ERROR

**Description**

This message occurs when starting or restarting the GPO Security Apply job fails.

**Corrective Action**

Address the specific reason for the issue, and then start the job by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

**Syslog Message**

Failed to start a GPO Security Apply job run for Vserver %s, reason: %s.

**Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

**gpo.frmwork.updateFailed****Severity**

ERROR

**Description**

This message occurs when updating group policy object (GPO) tables fails.

**Corrective Action**

Address the specific reason for the issue, and then update the tables by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

**Syslog Message**

GPO tables update failed for Vserver %s, reason: %s.

**Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.frmwork.updateRSOPFailed

### Severity

ERROR

### Description

This message occurs when updating group policy object (GPO) RSoP (Resultant Set of Policy) settings fails.

### Corrective Action

Address the specific reason for the issue, and then update the settings by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

GPO RSoP update failed for Vserver %s, reason: %s.

### Parameters

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.job events

### gpo.job.mainUpdateFailed

#### Severity

ERROR

#### Description

This message occurs when the GPO Main Update job run fails.

#### Corrective Action

Address the specific reason for the issue, and then run the job again by running the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

#### Syslog Message

GPO Main Update job (%s) run failed for Vserver %s, reason: %s.

#### Parameters

**jobid** (STRING): GPO job ID.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

### gpo.job.mainUpdateSuccess

#### Severity

INFORMATIONAL

#### Description

This message occurs when the GPO Main Update job run succeeds.

**Corrective Action**

(None).

**Syslog Message**

Successfully completed a GPO Main Update job (%s) run for Vserver %s.

**Parameters**

**jobid** (STRING): GPO job ID.

**vserver** (STRING): Vserver name.

**gpo.job.outOfMemory****Severity**

ERROR

**Description**

This event occurs when the system has not enough memory to perform the GPO (Group Policy Object) job.

**Corrective Action**

Ensure there is enough memory available to perform the GPO job.

**Syslog Message**

GPO job %s failure: Out of memory.

**Parameters**

**jobid** (STRING): The GPO job ID.

**gpo.job.securityApplyFailed****Severity**

ERROR

**Description**

This message occurs when the GPO Security Apply job run fails.

**Corrective Action**

Address the specific reason for the issue, and then run the job again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

**Syslog Message**

GPO Security Apply job (%s) run failed for Vserver %s, reason: %s.

**Parameters**

**jobid** (STRING): GPO job ID.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

**gpo.job.securityApplySuccess**

**Severity**

INFORMATIONAL

**Description**

This message occurs when the GPO Security Apply job run succeeds.

**Corrective Action**

(None).

**Syslog Message**

Successfully completed a GPO Security Apply job (%s) run for Vserver %s.

**Parameters**

**jobid** (STRING): GPO job ID.

**vserver** (STRING): Vserver name.

## gpo.reg events

### gpo.reg.applied

**Severity**

INFORMATIONAL

**Description**

This message occurs when the group policy object (GPO) registry policy settings are applied.

**Corrective Action**

(None).

**Syslog Message**

Successfully applied GPO registry settings to Vserver %s.

**Parameters**

**vserver** (STRING): Vserver name.

### gpo.reg.bCache.applied

**Severity**

INFORMATIONAL

**Description**

This message occurs when the group policy object (GPO) BranchCache settings are applied.

**Corrective Action**

(None).

**Syslog Message**

Successfully applied BranchCache settings to Vserver %s.

## Parameters

**vserver** (STRING): Vserver name.

## **gpo.reg.bCache.notApplied**

### Severity

ERROR

### Description

This message occurs when applying group policy object (GPO) registry BranchCache policy settings fails.

### Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

Failed to apply BranchCache settings for Vserver %s, reason: %s.

## Parameters

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## **gpo.reg.notApplied**

### Severity

ERROR

### Description

This message occurs when applying group policy object (GPO) registry policy settings fails.

### Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

Failed to apply GPO registry settings for Vserver %s, reason: %s.

## Parameters

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## **gpo.sec events**

### **gpo.sec.EvtAdt.applied**

### Severity

INFORMATIONAL

### Description

This message occurs when applying group policy object (GPO) event audit policy settings succeeds.

**Corrective Action**

(None).

**Syslog Message**

Successfully applied event audit and log settings to Vserver %s.

**Parameters**

**vserver** (STRING): Vserver name.

**gpo.sec.EvtAdt.notApplied****Severity**

ERROR

**Description**

This message occurs when applying group policy object (GPO) event audit policy settings fails.

**Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

**Syslog Message**

Failed to apply event audit and log settings to Vserver %s, reason: %s.

**Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

**gpo.sec.fsec.applied****Severity**

INFORMATIONAL

**Description**

This message occurs when the group policy object (GPO) file security settings are applied.

**Corrective Action**

(None).

**Syslog Message**

Successfully applied file security settings to Vserver %s.

**Parameters**

**vserver** (STRING): Vserver name.

**gpo.sec.fsec.notApplied****Severity**

ERROR

## Description

This message occurs when applying group policy object (GPO) file security settings fails.

## Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

## Syslog Message

Failed to apply file security settings to Vserver %s, reason: %s.

## Parameters

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.sec.fsec.notAppliedFor

### Severity

NOTICE

### Description

This message occurs when applying a group policy object (GPO) security descriptor to a given file fails.

### Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

Failed to apply file security settings to path %s in Vserver %s, reason: %s.

### Parameters

**path** (STRING): File and directory path to which a security descriptor is to be applied.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.sec.kerb.applied

### Severity

INFORMATIONAL

### Description

This message occurs when applying group policy object (GPO) Kerberos settings succeeds.

### Corrective Action

(None).

### Syslog Message

Successfully applied Kerberos settings to Vserver %s.

### Parameters

**vserver** (STRING): Vserver name.



## **gpo.sec.kerb.notApplied**

### **Severity**

ERROR

### **Description**

This message occurs when applying group policy object (GPO) Kerberos settings fails.

### **Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### **Syslog Message**

Failed to apply Kerberos settings to Vserver %s, reason: %s.

### **Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## **gpo.sec.privR.applied**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when applying group policy object (GPO) privilege-rights settings succeeds.

### **Corrective Action**

(None).

### **Syslog Message**

Successfully applied privilege rights settings to Vserver %s.

### **Parameters**

**vserver** (STRING): Vserver name.

## **gpo.sec.privR.notApplied**

### **Severity**

ERROR

### **Description**

This message occurs when applying group policy object (GPO) privilege-rights settings fails.

### **Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### **Syslog Message**

Failed to apply privilege-rights settings to Vserver %s, reason: %s.

## Parameters

**vserver** (STRING): Vserver name.  
**reason** (STRING): Reason for the failure.

## **gpo.sec.privR.notAppliedFor**

### Severity

NOTICE

### Description

This message occurs when applying group policy object (GPO) privilege-rights fails for a supported setting.

### Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

Failed to apply privilege-rights %s setting to Vserver %s, reason: %s.

## Parameters

**setting** (STRING): Name of the setting: Take ownership, Security privilege, or Bypass traverse.  
**vserver** (STRING): Vserver name.  
**reason** (STRING): Reason for the failure.

## **gpo.sec.reg.applied**

### Severity

INFORMATIONAL

### Description

This message occurs when applying group policy object (GPO) security registry settings succeeds.

### Corrective Action

(None).

### Syslog Message

Successfully applied security registry settings to Vserver %s.

## Parameters

**vserver** (STRING): Vserver name.

## **gpo.sec.reg.notAppliedFor**

### Severity

ERROR

### Description

This message occurs when applying a group policy object (GPO) security registry setting fails.

## Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

## Syslog Message

Failed to apply security registry %s setting to Vserver %s, reason: %s.

## Parameters

**setting** (STRING): Name of the setting: SMB signing.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## gpo.sec.resAno.applied

### Severity

INFORMATIONAL

### Description

This message occurs when applying group policy object (GPO) security restrict anonymous settings succeeds.

### Corrective Action

(None).

### Syslog Message

Successfully applied restrict anonymous settings to Vserver %s.

### Parameters

**vserver** (STRING): Vserver name.

## gpo.sec.resAno.notApplied

### Severity

ERROR

### Description

This message occurs when applying group policy object (GPO) security restrict anonymous settings fails.

### Corrective Action

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### Syslog Message

Failed to apply restrict anonymous settings to Vserver %s, reason: %s.

### Parameters

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## **gpo.sec.resGrp.applied**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when applying group policy object (GPO) restricted-groups settings succeeds.

### **Corrective Action**

(None).

### **Syslog Message**

Successfully applied restricted group settings to Vserver %s.

### **Parameters**

**vserver** (STRING): Vserver name.

## **gpo.sec.resGrp.notApplied**

### **Severity**

ERROR

### **Description**

This message occurs when applying group policy object (GPO) restricted-groups settings fails.

### **Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### **Syslog Message**

Failed to apply restricted-groups settings to Vserver %s, reason: %s.

### **Parameters**

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for the failure.

## **gpo.sec.resGrp.notAppliedFor**

### **Severity**

NOTICE

### **Description**

This message occurs when applying group policy object (GPO) restricted-groups settings fails for a given group.

### **Corrective Action**

Address the specific reason for the issue, and then apply the settings again by entering the "vserver cifs group-policy update -force-reapply-all-settings true -vserver <vserver-name>" command.

### **Syslog Message**

Failed to apply restricted-groups settings for group %s in Vserver %s, reason: %s.

## Parameters

**group** (STRING): Group name.  
**vserver** (STRING): Vserver name.  
**reason** (STRING): Reason for the failure.

# gpo.security events

## gpo.security.unrecogFileForm

### Severity

ERROR

### Description

This event occurs when the GptTmpl.inf file is not in UNICODE (16 bit)format.

### Corrective Action

Ensure the security file GptTmpl.inf is in UNICODE 16bit format.

### Syslog Message

Security file is not unicoded.

### Parameters

**filename** (STRING): The security filename.

## gpo.security.unrecogTemplate

### Severity

ERROR

### Description

This event occurs when the system fails to recognize the security template.

### Corrective Action

Ensure the security file GptTmpl.inf has the correct header.

### Syslog Message

Unable to recognize the security template.

### Parameters

(None).

# ha events

## ha.aggrmigrduration events

### ha.aggrMigrDuration

#### Severity

NOTICE

#### Description

This message occurs every 60 seconds during an aggregate migration operation. It indicates the current phase of the aggregate migration, the high-availability (HA) operation that triggered the aggregate migration, and the elapsed time since the start of the current phase of the aggregate migration.

#### Corrective Action

(None).

#### Syslog Message

As part of the current %s operation, the aggregate migration phase %s has been ongoing for %lld msec.  
Aggregate undergoing migration: %s.

#### Parameters

**aggr\_migr\_op\_name** (STRING): Name of the HA operation that triggered the aggregate migration.

**aggr\_migr\_phase** (STRING): Current phase of the aggregate migration.

**elapsed\_time** (LONGINT): Elapsed time since the start of the current phase of the aggregate migration.

**aggr\_name** (STRING): Name of the aggregate undergoing migration.

## ha.bdfu events

### ha.bdfu.notEnabledOnDest

#### Severity

INFORMATIONAL

#### Description

This message occurs during storage failover (SFO) giveback or aggregate relocation when the source node cannot reenables background disk firmware update (BDFU) on the destination node. BDFU will be reenables automatically on the destination node after all the aggregates have been given back.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**nodename** (STRING): Destination node.

# ha.disasterbootinitiated events

## ha.disasterBootInitiated

### Severity

NOTICE

### Description

This message occurs when an HA disaster recovery boot is initiated. Such a boot occurs when the HA partner of a failing node also experiences a failure, and one from which it cannot recover.

### Corrective Action

(None).

### Syslog Message

HA disaster recovery boot initiated.

### Parameters

(None).

# ha.fm events

## ha.fm.platformTypeChanged

### Severity

NOTICE

### Description

This message occurs when the Failover Monitor (FM) platform type changes in response to changes in the cluster configuration.

### Corrective Action

(None).

### Syslog Message

FM platform type changed from %s to %s.

### Parameters

**from** (STRING): Previous FM platform type.

**to** (STRING): New FM platform type.

# ha.gb events

## ha.gb.mbxAccessErr

### Severity

ALERT

## Description

This message occurs when a mailbox disk cannot be accessed prior to a giveback.

## Corrective Action

Verify the mailbox disk physical locations and mailbox disk counts by using the "storage failover mailbox-disk show -node node-id" advanced command. If all mailbox disks are present for the indicated physical location, retry the "storage failover giveback" command or sf-send-home-initiate ZAPI. If none of the mailbox disks for the indicated physical location are shown, wait a few minutes, and then retry the show command to see whether all mailbox disks are present. If all mailbox disks are present, retry the "storage failover giveback" command or sf-send-home-initiate ZAPI. If some but not all mailbox disks are present for the indicated physical location, then set the "-override-vetoes" field to "true" for either the "storage failover giveback" command or sf-send-home-initiate ZAPI, and then retry the giveback.

## Syslog Message

Problem accessing %s mailbox disks prior to giveback. Found %d disks, expected %d disks. Override vetoes option: %s; Internal error code: %d status: %d.

## Parameters

**physical\_location** (STRING): Identifies the physical location of the mailbox disks. Values are Local Node, Partner Node, Mediator, or a node name.

**mailbox\_disk\_count** (INT): Current count of mailbox disks having the physical location specified by the physical\_location parameter.

**expected\_count** (INT): Expected count of mailbox disks having the physical location specified by the physical\_location parameter.

**override\_vetoes** (STRING): Value of the "-override-vetoes" field in the "storage failover giveback" command or sf-send-home-initiate ZAPI.

**error\_code** (INT): Disk I/O or other error code.

**error\_status** (INT): Additional internal status related to the error code.

## ha.gb.mbxCountMismatch

### Severity

NOTICE

### Description

This message occurs when a mailbox disk cannot be accessed prior to a giveback, but giveback is allowed to proceed because the "override-vetoes" field is set to "true" for the "storage failover giveback" command or the sf-send-home-initiate ZAPI.

### Corrective Action

(None).

### Syslog Message

Problem accessing %s mailbox disks during giveback. Found %d disks, expected %d disks. Override vetoes option: %s; Internal result code: %d status: %d"

### Parameters

**physical\_location** (STRING): Identifies the physical location of the mailbox disks. Values are Local Node, Partner Node, Mediator, or a node name.

**mailbox\_disk\_count** (INT): Current count of mailbox disks having the physical location specified by the physical\_location parameter.

**expected\_count** (INT): Expected count of mailbox disks having the physical location specified by the physical\_location parameter.



**override\_vetoes** (STRING): Value of the "-override-vetoes" field in the "storage failover giveback" command or sf-send-home-initiate ZAPI.

**result\_code** (INT): Disk I/O or other result code.

**result\_status** (INT): Additional internal status related to the result code.

## ha.gb.mbxNotFound

### Severity

ALERT

### Description

This message occurs when a mailbox disk cannot be accessed prior to a giveback.

### Corrective Action

Verify the mailbox disk physical locations and mailbox disk counts by using the "storage failover mailbox-disk show -node node-id" advanced command. If all mailbox disks are present, retry the "storage failover giveback" command or the sf-send-home-initiate ZAPI. If none of the mailbox disks for the indicated mailbox owner are shown, wait a few minutes, and then retry the show command to see whether all mailbox disks are present. If all mailbox disks are present, retry the "storage failover giveback" command or sf-send-home-initiate ZAPI. If some but not all mailbox disks are present for the indicated mailbox owner, verify whether all nodes in the cluster are healthy and accessible by using the "cluster show" command, and check whether disks are healthy by using the "disk show" command.

### Syslog Message

Problem accessing %s mailbox disks prior to giveback. Found %d disks. Override vetoes option: %s; Internal error code: %d status: %d.

### Parameters

**mailbox\_owner** (STRING): Identifies the node in the high-availability (HA) pair that owns the missing mailbox disk.

**mailbox\_disk\_count** (INT): Number of mailbox disks found. These disks are owned by the node specified by the mailbox\_owner parameter.

**override\_vetoes** (STRING): Value of the "-override-vetoes" parameter passed to the "storage failover giveback" command or the sf-send-home-initiate ZAPI.

**error\_code** (INT): Disk I/O or other error code.

**error\_status** (INT): Additional internal status related to the error code.

## ha.get takeoverlockstatus events

### ha.getT takeoverLockStatus

### Severity

NOTICE

### Description

This message occurs when the HA subsystem attempts to acquire takeover lock as part of the takeover operation.

### Corrective Action

(None).

## Syslog Message

Attempting to acquire takeover lock. Status: %s.

## Parameters

**status** (STRING): Status of the attempt to acquire the lock.

# ha.giveback events

## ha.giveback.aggrOnlineFailed

### Severity

INFORMATIONAL

### Description

This message occurs when the RAID or WAFL® subsystem fails to bring online the aggregate or its containing volumes after giveback.

### Corrective Action

Check the EMS logs for messages from the RAID and WAFL subsystems to determine the exact cause of the failure and corrective action.

## Syslog Message

Bringing online of aggregate %s or its volumes failed after giveback with error %d.

## Parameters

**aggrName** (STRING): Name of the aggregate that was given back.

**error** (INT): Error that occurred while bringing online the aggregate or its volumes after giveback.

## ha.giveback.aggrOnlineTimeout

### Severity

ERROR

### Description

This message occurs when a node fails to bring online the aggregate and its containing volumes within the specified time after giveback.

### Corrective Action

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online.

## Syslog Message

Node did not bring online the aggregate %s and its volumes after the giveback for %llu milliseconds.

## Parameters

**aggrUUID** (STRING): UUID of the aggregate that was given back.

**time** (LONGINT): Timeout, in milliseconds, to bring online the aggregate and its volumes after the giveback.

## ha.giveback.aggrPartnerTimeout

**Severity**

INFORMATIONAL

**Description**

This message occurs during a giveback operation, when a node does not receive notification from its high-availability (HA) partner within the specified time that the aggregate is online.

**Corrective Action**

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online on the partner node.

**Syslog Message**

Node did not receive notification from its partner in %llu milliseconds that aggregate %s is online after the giveback.

**Parameters**

**time** (LONGINT): Timeout, in milliseconds, for notification from the partner node that an aggregate is online after giveback.

**aggrName** (STRING): Name of the aggregate that was given back.

**ha.giveback.commWithPartnerFailed****Severity**

ERROR

**Description**

This message occurs when the local node cannot communicate over the cluster network with the d-blade of its high-availability (HA) partner while giving back an aggregate.

**Corrective Action**

1. Verify that the partner node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local and the partner node are up.
3. If the preceding checks pass, then run the "storage failover giveback" command to initiate giveback. To force giveback, use the "storage failover giveback -require-partner-waiting false" command.

**Syslog Message**

Failed to communicate with the partner %s over the cluster network while giving back aggregate %s.

**Parameters**

**partnerName** (STRING): Name of the partner node.

**aggrName** (STRING): Name of the aggregate being given back.

**ha.giveback.connectToPartnerFailed****Severity**

ERROR

**Description**

This message occurs when the local node cannot connect over the cluster network to the d-blade of its high-availability (HA) partner while giving back an aggregate.

## Corrective Action

1. Verify that the partner node is up. 2. Use the "network interface show" command to verify that the cluster network interfaces on the local node and the partner node are up. 3. If the preceding checks pass, then run the "storage failover giveback" command to initiate giveback. To force giveback, use the "storage failover giveback -require-partner-waiting false" command.

## Syslog Message

Failed to connect with the partner node %s (d-blade ID: %s) over the cluster network while giving back aggregate %s.

## Parameters

- partnerName** (STRING): Name of the partner node.
- partnerDbladeid** (STRING): D-blade ID of the partner node.
- aggrName** (STRING): Name of the aggregate being given back.

## ha.giveback.emptyAggr

### Severity

ERROR

### Description

This message occurs when the local node does not find any disks in the aggregate being given back.

### Corrective Action

1. Use the "storage disk show -aggregate aggrName" command to verify that the aggregate has disks. 2. If the preceding check passes, then run the "storage failover giveback" command to initiate giveback.

### Syslog Message

Aggregate %s does not have any disks.

### Parameters

- aggrName** (STRING): Name of the aggregate being given back.

## ha.giveback.sysCommit

### Severity

INFORMATIONAL

### Description

This message occurs during the commit phase of the giveback of an aggregate to indicate the time taken by the specified subsystem to commit.

### Corrective Action

(None).

### Syslog Message

Subsystem %s took %llu msecs to commit giveback of aggregate '%s'.

### Parameters

- subsys** (STRING): Name of the subsystem performing the commit.
- time** (LONGINT): Time taken by the subsystem to commit the giveback.

**aggr** (STRING): Name of the aggregate being given back.

## ha.giveback.sysVetoChk

### Severity

INFORMATIONAL

### Description

This message occurs during the veto phase of the giveback of an aggregate to indicate the time taken by the specified subsystem to check whether giveback should be aborted.

### Corrective Action

(None).

### Syslog Message

Subsystem %s took %llu msec to check whether to veto giveback of aggregate '%s'.

### Parameters

**subsys** (STRING): Name of the subsystem performing the veto check.

**time** (LONGINT): Time taken by the subsystem to check whether giveback should be aborted.

**aggr** (STRING): Name of the aggregate being given back.

## ha.giveback.totalCommit

### Severity

INFORMATIONAL

### Description

This message occurs during the giveback of an aggregate to indicate the total time taken by the commit phase.

### Corrective Action

(None).

### Syslog Message

Total time taken to commit the giveback of aggregate '%s' was %llu msec.

### Parameters

**aggr** (STRING): Name of the aggregate being given back.

**time** (LONGINT): Total time taken to commit the giveback.

## ha.giveback.usingHAInterconnect

### Severity

INFORMATIONAL

### Description

This message occurs when the local node uses the high-availability (HA) interconnect instead of the preferred cluster network to communicate with its HA partner during giveback.

**Corrective Action**

(None).

**Syslog Message**

Local node used the HA interconnect to communicate with the partner during giveback of aggregate %s.

**Parameters**

**aggrName** (STRING): Name of the aggregate being given back.

## ha.localnodeshutdown events

**ha.localNodeShutDown****Severity**

NOTICE

**Description**

This message occurs when shutdown of the local node is initiated.

**Corrective Action**

(None).

**Syslog Message**

Shutdown of the local node has been initiated with inhibit\_takeover set to %s.

**Parameters**

**inhibit\_takeover** (STRING): Specifies whether takeover by the partner node is inhibited on shutdown of the local node.

## ha.mbxresvreleasedelay events

**ha.mbxResvReleaseDelay****Severity**

NOTICE

**Description**

This message occurs when the takeover node cannot release a disk reservation for a mailbox disk in preparation for a giveback event. The mailbox reservation release will be automatically retried after a few minutes.

**Corrective Action**

The mailbox reservation release will be automatically retried after a few minutes. If this problem persists and "storage failover giveback" is not run automatically or cannot be run manually, release the mailbox disks reservations manually using maintenance mode.

**Syslog Message**

Problem releasing reservation for mailbox disk attached to node with System ID %d. Disk name %s, disk owner %s, disk UUID: %s, status: %d.

## Parameters

**physical\_location** (LONGINT): System ID of the node directly attached to the mailbox disk on which the reservation release failure occurred.

**disk\_name** (STRING): Name of the disk for which the reservation release failed.

**mbx\_disk\_owner** (STRING): Owner of the mailbox disk where the reservation release failed. Values are Local Node or Partner Node.

**disk\_uid** (STRING): Unique identifier of the disk where the reservation release failed.

**status** (INT): Reservation release status code.

# ha.netpartition events

## ha.netPartition

### Severity

NOTICE

### Description

This message occurs when a Data ONTAP® subsystem detects an error potentially due to network partitioning, but cannot take over the partner node or allow the partner to take over because both takeover-of-partner and takeover-by-partner are both disabled. Network partitioning is indicated by concurrent errors on more than one of a node's interfaces, if a halt, reboot, takeover or similar action affecting connectivity is not already in progress.

### Corrective Action

(None).

### Syslog Message

Network partition error. Internal error code %d; status: 0x%x; internal request ID: %d.

## Parameters

**error\_code** (INT): Internal error code.

**error\_status** (INTEX): Additional internal status bits related to error\_code.

**req\_id** (INT): Internal request ID.

**duration** (LONGINT): Elapsed time in msec spent resolving the reported error.

## ha.netPartition.reboot

### Severity

NOTICE

### Description

This message occurs when a Data ONTAP® subsystem detects an error potentially due to network partitioning, and reboots the local node due to the network partition. The reboot occurs only if ` takeover-by-partner is enabled on the local node. Network partitioning is indicated by concurrent errors on more than one of a node's interfaces, if a halt, reboot, takeover or similar action affecting connectivity is not already in progress.

### Corrective Action

(None).

## Syslog Message

Rebooting the local node due to a network partition error. Internal error code %d; status: 0x%x; internal request ID: %d.

## Parameters

**error\_code** (INT): Internal error code.

**error\_status** (INHEX): Additional internal status bits related to error\_code.

**req\_id** (INT): Internal request ID.

**duration** (LONGINT): Elapsed time in msec spent resolving the reported error.

## ha.netPartition.takeover

### Severity

NOTICE

### Description

This message occurs when a Data ONTAP® subsystem detects an error potentially due to network partitioning, and initiates a takeover of the partner node due to the partition. Network partitioning is indicated by concurrent errors on more than one of a node's interfaces, if a halt, reboot, takeover or similar action affecting connectivity is not already in progress.

### Corrective Action

(None).

## Syslog Message

Taking over the partner node due to a network partition error. Internal error code %d; status: 0x%x; internal request ID: %d.

## Parameters

**error\_code** (INT): Internal error code.

**error\_status** (INHEX): Additional internal status bits related to error\_code.

**req\_id** (INT): Internal request ID.

**duration** (LONGINT): Elapsed time in msec spent resolving the reported error.

## ha.netparttodisabled events

### ha.netPartToDisabled

### Severity

NOTICE

### Description

This message occurs when takeovers due to network partitioning are disabled because of the number of previous, consecutive, network-partition takeovers that occurred within a preceding detection period. Subsequent takeovers due to network partitioning are automatically reenabled after the disable time limit is reached or after a takeover is initiated for a different reason.

### Corrective Action

(None).



## Syslog Message

Network partitioning takeovers are disabled due to %d network-related takeovers in the preceding %d minutes. Takeovers by local node %d; partner node %d. Takeovers will be re-enabled in %d minutes or when another type of takeover occurs.

## Parameters

**takeovers\_initiated** (INT): Consecutive takeovers initiated from either node within the detection period that were due to network partitioning.

**detection\_period** (INT): Number of minutes over which the network partition takeover attempts were monitored and counted.

**local\_node\_initiated** (INT): Number of takeovers due to network partitioning occurring within the detection period that were initiated by the local node. This number is included in the takeover\_initiated number.

**partner\_node\_initiated** (INT): Number of takeovers due to network partitioning occurring within the detection period that were initiated by the partner node. This number is included in the takeover\_initiated number.

**minutes\_until\_reenable** (INT): Minutes until takeover due to network partitioning is reenabled.

# ha.netra events

## ha.netra.dskChkNoContDsk

### Severity

ERROR

### Description

This message occurs during an aggregate migration, when the destination node either cannot determine whether at least one disk in the aggregate is a disk partition, or the associated container disk is not visible to the destination node. The aggregate migration operation is aborted.

### Corrective Action

Use the (privilege: advanced) "storage disk refresh-ownership" command to rescan the disks and update ownership information. After refreshing the disk ownership info, use the "storage show disk" command to view the ownership of all disks, and the "storage show disk -container-type shared" command to view container-disk ownership. If any disks are incorrectly owned, update the disk ownership information. If the problem persists, use the "storage failover show -fields local-missing-disks, partner-missing-disks" command to check whether there are any disks that are not visible to a subset of the nodes. Use the "storage failover show" command to check whether the high-availability (HA) partner's information is known. If HA partner information is missing, verify the HA Interconnect between both nodes is installed and functioning correctly. If not, fix the HA Interconnect cabling issues to resolve the connectivity problems. Use the "storage disk show" command to check whether there are any failed disks. If failed disks are present, remove or replace them. If there are no failed disks, check the cabling for loose connections or a bad cable on the node where the disk is not visible. If the problem persists, contact NetApp technical support.

## Syslog Message

Disk with UUID %s is not visible to node %s, which is the destination of an aggregate migration operation. As a result, the aggregate migration was aborted. Check the event log of the destination node to determine the name of the aggregate whose migration was aborted.

## Parameters

**disk\_uuid** (STRING): Universal unique identifier (UUID) of the disk that is not visible to the destination node.

**node\_name** (STRING): Name of the destination node.

## ha.nomediatormbx events

### ha.noMediatorMbx

#### Severity

ERROR

#### Description

This message occurs when the non-local mailbox disks of a node are not accessible. The system will not be able to perform a takeover and can potentially prevent a takeover by the partner as well.

#### Corrective Action

Use the "storage failover show -fields local-mailbox-disks" command to verify that all nodes in the cluster can access their respective mailbox disks. Use the (privilege: advanced) "storage iscsi-initiator show" command to verify that all iSCSI connections are up and healthy. If not healthy, use the "storage iscsi-initiator add-target" and "storage iscsi-initiator connect" commands to add and connect any missing or misconnected iSCSI targets.

#### Syslog Message

Takeover not possible because non-local mailbox disks owned by node "%s" cannot be accessed.

#### Parameters

`mailbox_owner` (STRING): Node that owns the non-local mailbox disks that could not be accessed.

## ha.noprogrammaticto events

### ha.noProgrammaticTO

#### Severity

NOTICE

#### Description

This message occurs when a takeover is initiated programmatically by Data ONTAP® and the takeover request fails because giveback or aggregate relocation cleanup is still in progress.

#### Corrective Action

(None).

#### Syslog Message

Failover monitor: Programmatic takeover failed due to giveback or aggregate relocation post-processing cleanup not yet complete.

#### Parameters

(None).

## ha.panicinfonotsent events

### ha.panicInfoNotSent

## Severity

NOTICE

## Description

This message occurs when a node cannot send panic information to its high-availability (HA) partner over the HA interconnect. The HA partner might still initiate takeover if it detects that the partner is down by other means, and takeover of the panicking node is allowed.

## Corrective Action

(None).

## Syslog Message

Local node cannot send a panic information message to the HA partner over the HA interconnect due to error: %s. Partner name: %s. Partner system ID: %llu.

## Parameters

**error** (STRING): Reason the node cannot send panic information to its HA partner.

**partner\_name** (STRING): Name of the HA partner.

**partner\_systemid** (LONGINT): System ID of the HA partner.

# ha.panicinfonotsentic events

## ha.panicInfoNotSentIC

## Severity

NOTICE

## Description

This message occurs when a node cannot send panic information to its high-availability (HA) partner over the HA interconnect due to an HA interconnect error. The HA partner might still initiate takeover if it detects that the partner is down by other means, and takeover of the panicking node is allowed.

## Corrective Action

(None).

## Syslog Message

Local node cannot send a panic information message to its HA partner due to an HA interconnect error: %s (error code: %d). Partner name: %s. Partner system ID: %llu.

## Parameters

**error** (STRING): Reason the node cannot send panic information to its HA partner.

**error\_code** (INT): Specific error generated by the HA Interconnect. Example error codes: 1

(ERR\_ICDOWN): Interconnect is down 4 (ERR\_SYSERROR): System error 5 (ERR\_ICERROR):

Interconnect error 6 (ERR\_OUTOFMEMORY): Out of memory 11 (ERR\_SENDSIZE): Send size is too large

19 (ERR\_CKSUM): Checksum error

**partner\_name** (STRING): Name of the HA partner.

**partner\_systemid** (LONGINT): System ID of the HA partner.

# ha.panicinforcvbadvrsn events

## ha.panicInfoRcvBadVrsn

### Severity

NOTICE

### Description

This message occurs when the version of the panic information received from the high-availability (HA) partner is not equal to the expected version. The receiving node will discard this panic information message. A takeover might still occur if the local node detects that its HA partner is down by other means, and takeover of the panicking node is allowed.

### Corrective Action

(None).

### Syslog Message

Panic information message received from the HA partner is the wrong version. Expected version: %u. Received version: %u. Partner name: %s. Partner system ID: %llu.

### Parameters

**expected\_version** (LONGINT): Expected version of the panic information.

**received\_version** (LONGINT): Received version of the panic information.

**partner\_name** (STRING): Name of the HA partner.

**partner\_systemid** (LONGINT): System ID of the HA partner.

# ha.panicinfosent events

## ha.panicInfoSent

### Severity

NOTICE

### Description

This message occurs when a node successfully sends panic information to its high-availability (HA) partner over the HA interconnect.

### Corrective Action

(None).

### Syslog Message

Node successfully sent a panic information message to its HA partner. Partner name: %s. Partner system ID: %llu.

### Parameters

**local\_hostname** (STRING): Name of the node that panicked and sent panic information over the HA interconnect.

**local\_systemid** (LONGINT): System ID of the node that panicked and sent panic information over the HA interconnect.

# ha.panicinfosznotok events

## ha.panicInfoSzNotOk

### Severity

NOTICE

### Description

This message occurs when the size of the panic information received from the high-availability (HA) partner is not equal to the expected panic information size. The receiving node will discard this panic information message. A takeover might still occur if the local node detects that its HA partner is down by other means, and takeover of the panicking node is allowed.

### Corrective Action

(None).

### Syslog Message

Panic information message received from the HA partner is the wrong size. Expected size: %u. Received size: %u. Partner name: %s. Partner system ID: %llu.

### Parameters

**expected\_size** (LONGINT): Expected size of the panic information.

**received\_size** (LONGINT): Received size of the panic information.

**partner\_name** (STRING): Name of the HA partner.

**partner\_systemid** (LONGINT): System ID of the HA partner.

# ha.resvconflicthalt events

## ha.resvConflictHalt

### Severity

NOTICE

### Description

This message occurs when the node halts itself after it finds disks with reservation conflicts. Reservations are placed on disks as part of a node taking over its partner. This disk reservation conflict is due to this node being taken over by its partner.

### Corrective Action

(None).

### Syslog Message

A disk reservation conflict was detected on disk %s at %-20s. Typically, this only occurs when the node was taken over by its partner.

### Parameters

**disk\_name** (STRING): Name of the disk.

**time** (STRING): Time when the disk reservation was detected.

**last\_update** (INT): Time passed since last rescan, in milliseconds.

# ha.resvmediatormbxstatus events

## ha.ResvMediatorMbxStatus

### Severity

NOTICE

### Description

This message occurs when the local node attempts to reserve a mailbox disk while trying to acquire the takeover lock as part of the takeover operation.

### Corrective Action

(None).

### Syslog Message

Attempting to reserve mailbox disk %s (UUID: %s) on node "%s" as part of takeover lock. Status: %d.

### Parameters

**mailbox\_disk\_name** (STRING): Name of the disk being reserved.

**mailbox\_disk\_uid** (STRING): UUID of the mailbox disk on which a reservation was attempted.

**mailbox\_host** (STRING): Name of the node that owns the mailbox disk.

**mailbox\_op\_status** (INT): Result of the disk reservation attempt on the mailbox disk.

# ha.takeover events

## ha.takeover.bypOptimChange

### Severity

INFORMATIONAL

### Description

This message occurs when the value of the 'bypass-takeover-optimization' parameter of the "storage failover" command is changed on a node.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**node\_name** (STRING): Name of the node.

**new\_value** (STRING): New value of the 'bypass-takeover-optimization' parameter.

## ha.takeover.checkEnd

### Severity

ALERT

**Description**

This message occurs when a node is simulating a takeover. This marks the end of a simulated takeover on the local node.

**Corrective Action**

(None).

**Syslog Message**

This node is testing the takeover workflow. This test is a simulation only. The simulation has ended on the local node.

**Parameters**

(None).

**ha.takeover.checkError****Severity**

ALERT

**Description**

This message occurs when a takeover with the check-only flag detects an error that would prevent takeover.

**Corrective Action**

Use the "storage failover check-takeover" command to identify the failures that occurred along the takeover path. Perform corrective action based on error information shown by the command.

**Syslog Message**

This node is testing the takeover workflow. This test is a simulation only. The simulation has encountered an error that would disrupt takeover.

**Parameters**

(None).

**ha.takeover.checkStart****Severity**

ALERT

**Description**

This message occurs when a node is simulating a takeover. This marks the start of a simulated takeover on the local node.

**Corrective Action**

(None).

**Syslog Message**

This node is testing the takeover workflow. This test is a simulation only. The simulation has started on the local node.

## Parameters

(None).

# ha.takeoverimpdegraded events

## ha.takeoverImpDegraded

### Severity

ALERT

### Description

This message occurs when the partner node is down and it is currently impossible to perform a takeover because the partner mailbox disks are not available.

### Corrective Action

Check connectivity to the partner disks. Identify the mailbox disks on the partner node by using the "storage failover show -instance" command. Verify that this node has connectivity to the partner mailbox disks by using the "storage disk show -instance" command. Resolve any connectivity issues by fixing cabling or shelf issues.

### Syslog Message

Takeover of the partner node is impossible due to lack of connectivity to the partner mailbox disks.

## Parameters

(None).

# ha.takeoverimphotshelf events

## ha.takeoverImpHotShelf

### Severity

EMERGENCY

### Description

This message occurs when the partner node is down and it is currently impossible to perform a takeover due to disk shelves that are overheating.

### Corrective Action

Check cooling for the disk shelves.

### Syslog Message

Takeover of the partner node is impossible due to disk shelves that are overheating.

## Parameters

(None).

# ha.takeoverimpic events



## ha.takeoverImpIC

### Severity

ALERT

### Description

This message occurs when the partner node is down and the local node cannot perform a takeover because of interconnect errors.

### Corrective Action

Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenable the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. When one link is down on platforms with two available links, HA link redundancy is lost. However, HA services are still working and enabled on the secondary link.

### Syslog Message

Takeover of the partner node is not possible due to interconnect errors.

### Parameters

(None).

## ha.takeoverimplowmem events

### ha.takeoverImpLowMem

### Severity

ERROR

### Description

This message occurs when the partner node is down; however, due to this node being low on memory, it is currently impossible to perform a takeover.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Takeover of the partner node is impossible due to a low memory situation.

### Parameters

(None).

## ha.takeoverimpnotdef events

### ha.takeoverImpNotDef

### Severity

ERROR

## Description

This message occurs when the partner node is down and it is currently impossible to perform a takeover.

## Corrective Action

Check the reason string and repair the problem reported.

## Syslog Message

Takeover of the partner node is impossible due to reason %s.

## Parameters

**reason** (STRING): Reason the takeover is currently impossible.

# ha.takeoverimpunsync events

## ha.takeoverImpUnsync

### Severity

ALERT

### Description

This message occurs when the partner node is down and it is currently impossible to perform a takeover because the takeover node does not have partner NVRAM data.

### Corrective Action

Verify that the interconnect cables are correctly connected between the nodes. Boot the partner node to reestablish the interconnect if it was shut down. If necessary, force a takeover with the '[run node] cf forcetakeover' command. This action might cause data loss if there is NVRAM data on the partner node not written to disk.

### Syslog Message

Takeover of the partner node is impossible due to lack of partner NVRAM data.

### Parameters

(None).

# ha.takeoverimpversion events

## ha.takeoverImpVersion

### Severity

ALERT

### Description

This message occurs when the partner node is down; however, due to a version mismatch, it is currently impossible to perform a takeover.

### Corrective Action

Upgrade the downrev node, so that both nodes are on the same release of Data ONTAP®. If a revert is being performed, downgrade the uprev node to the same release as the partner node.

## Syslog Message

Takeover of the partner node is impossible due to version mismatch.

## Parameters

(None).

# ha.tkover events

## ha.tkover.waitARLabort

### Severity

INFORMATIONAL

### Description

This message occurs when a pending takeover is waiting for an in-progress aggregate relocation (ARL) operation to complete or abort before proceeding with the takeover.

### Corrective Action

(None).

## Syslog Message

Node %s has been waiting for an ongoing ARL for %d seconds. It will proceed with takeover when the ARL operation has been completed or aborted.

## Parameters

**node\_name** (STRING): Name of the node doing the takeover.

**wait\_time** (LONGINT): Amount of time spent waiting for the ongoing ARL operation to be completed or aborted.

## ha.tkover.waitGBabort

### Severity

INFORMATIONAL

### Description

This message occurs when a pending takeover is waiting for an in-progress SFO giveback operation to complete or abort before proceeding with the takeover.

### Corrective Action

(None).

## Syslog Message

Node %s has been waiting for an ongoing SFO giveback for %d seconds. It will proceed with takeover when the SFO giveback operation has been completed or aborted.

## Parameters

**node\_name** (STRING): Name of the node doing the takeover.

**wait\_time** (LONGINT): Amount of time spent waiting for the ongoing SFO giveback operation to be completed or aborted.

# ha.tkoverabortednomediatormbx events

## ha.tkoverAbortedNoMediatorMbx

### Severity

NOTICE

### Description

This message occurs when a planned failover is aborted because the mediator mailbox disks were not detected.

### Corrective Action

Run the "storage failover show -fields local-mailbox-disks partner-mailbox-disks" and "storage failover show -instance" commands to check whether the mediator mailbox disks are missing. If the mediator mailbox disks are not present, rescan the disks on both nodes in the HA pair to determine the latest ownership by using the "storage disk refresh-ownership" command. If the rescan does not work, contact NetApp technical support.

### Syslog Message

(None).

### Parameters

(None).

# ha.toafterrcvpanicinfo events

## ha.toAfterRcvPanicInfo

### Severity

NOTICE

### Description

This message occurs when the local node initiates a takeover of its high-availability (HA) partner after detecting a panic of the HA partner node.

### Corrective Action

(None).

### Syslog Message

Local node initiated takeover after receiving a panic information message from its HA partner. Partner name: %s. Partner system ID: %llu.

### Parameters

**hostname** (STRING): Host name of the node being taken over.  
**systemid** (LONGINT): System ID of the node being taken over.

# ha.tp events

## ha.tp.testPointHit

### Severity

NOTICE

### Description

This message occurs when the testpoint is encountered or hit.

### Corrective Action

(None).

### Syslog Message

Testpoint hit: '%s' [%s, %s], [%s:%d], count %d ⇒ '%s'.

### Parameters

**message** (STRING): Test message.

**tpModuleName** (STRING): Testpoint module name.

**tpId** (STRING): ID of the testpoint within the module.

**fileName** (STRING): File name.

**lineNo** (INT): Line number.

**tpCount** (INT): Number of times the testpoint triggers before being disabled.

**tpType** (STRING): Type of testpoint.

# hamsg events

## hamsg.diskcheckresp events

### hamsg.diskCheckResp

#### Severity

NOTICE

#### Description

This message occurs when a node receives a disk inventory check response from its high-availability (HA) partner.

#### Corrective Action

(None).

#### Syslog Message

The local node has received a disk inventory check response (%d) from its HA partner.

#### Parameters

**response** (INT): The disk inventory check response received from the HA partner.

## hamsg.dskchkpopulateerr events

### hamsg.dskChkPopulateErr

#### Severity

ERROR

#### Description

This message occurs when a disk inventory check cannot be sent to the local node's high-availability (HA) partner because the disk check request could not be correctly constructed.

#### Corrective Action

Retry the operation. If the problem persists, use the "storage failover show -fields local-missing-disks, partner-missing-disks" command to rescan the disks and to determine the latest ownership. If the disk inventory mismatch is not resolved, check the cabling for either loose connections or for a bad cable on the partner node.

#### Syslog Message

A disk inventory check cannot be sent to the local node's HA partner. Failure reason: %s.

#### Parameters

**failure\_reason** (STRING): The reason that the disk inventory check request could not be properly constructed.

# haosc events

## haosc.config events

### haosc.config.unknown.cf

#### Severity

ALERT

#### Description

This message occurs when the chassis and controller field-replaceable unit programmable read-only memories (FRU PROMs) have inconsistent data.

#### Corrective Action

Boot into Maintenance mode and run the 'ha-config modify' command to set the controller and chassis configuration to stand-alone or HA, as appropriate. Exercise extreme caution. Setting the wrong configuration might lead to data loss.

#### Syslog Message

Cannot determine whether configuration should be stand-alone or HA. Chassis is in %s configuration, and controller is in %s configuration.

#### Parameters

**cha\_config** (STRING): Chassis FRU PROM HA configuration string.  
**cont\_config** (STRING): Controller FRU PROM HA configuration string.

### haosc.config.unknown.ls

#### Severity

ALERT

#### Description

This message occurs when the chassis and controller field-replaceable unit programmable read-only memory (FRU PROM) configurations conflict with the HA mode.

#### Corrective Action

Boot into Maintenance mode and run the 'ha-config modify' command to set the controller and chassis configuration to stand-alone or HA, as appropriate. Setting the wrong configuration might lead to data loss. If you need assistance, contact NetApp technical support.

#### Syslog Message

Cannot determine whether the configuration should be stand-alone or HA. Chassis is in %s configuration, controller is in %s configuration, and the HA mode is %s.

#### Parameters

**cha\_config** (STRING): Chassis FRU PROM HA configuration string.  
**cont\_config** (STRING): Controller FRU PROM HA configuration string.  
**mode** (STRING): HA mode string.

## haosc.config.unknown.nv

### Severity

ALERT

### Description

This message occurs when the chassis and controller field-replaceable unit programmable read-only memories (FRU PROMs) have factory default configurations and the nonvolatile memory is dirty.

### Corrective Action

Boot into Maintenance mode and run the 'ha-config modify' command to set the controller and chassis configuration to stand-alone or HA, as appropriate. Setting the wrong configuration might lead to data loss. If you need assistance, contact NetApp technical support.

### Syslog Message

Cannot determine whether configuration should be stand-alone or HA. Chassis is in %s configuration, controller is in %s configuration, and the nonvolatile memory is dirty.

### Parameters

**cha\_config** (STRING): Chassis FRU PROM HA configuration string.  
**cont\_config** (STRING): Controller FRU PROM HA configuration string.

## haosc.fru events

### haosc.fru.op.failed

### Severity

ALERT

### Description

This message occurs when the chassis or controller field-replaceable unit programmable read-only memory (FRU PROM) read/write fails.

### Corrective Action

Replace the system chassis or controller, and contact NetApp technical support for assistance.

### Syslog Message

%s FRU PROM %s operation failed.

### Parameters

**fru** (STRING): Chassis/controller string.  
**op** (STRING): Read/write operation string.

## haosc.invalid events

### haosc.invalid.config

### Severity

EMERGENCY



## Description

This message occurs when the chassis or controller field-replaceable unit programmable read-only memory (FRU PROM) returns an invalid configuration.

## Corrective Action

Boot into Maintenance mode and run the 'ha-config modify' command to set the controller and chassis configuration to stand-alone or HA, as appropriate. Setting the wrong configuration might lead to data loss. If you need assistance, contact NetApp technical support.

## Syslog Message

%s FRU PROM returns an invalid configuration.

## Parameters

fru (STRING): Chassis/controller string.

# haosc.sa events

## haosc.sa.config

### Severity

ERROR

### Description

This message occurs when nonvolatile memory on a stand-alone system is configured for a high-availability system. Only half of the available nonvolatile memory will be used and system performance will be affected.

### Corrective Action

Reboot the system to use all of the nonvolatile memory.

### Syslog Message

Only half of the available nonvolatile memory is in use. This will affect system performance.

### Parameters

(None).

# hashd events

## hashd.maxstore events

### hashd.maxstore.oversize

#### Severity

NOTICE

#### Description

This message occurs when the size of hashes in the hash store is greater than the maximum size configured.

#### Corrective Action

If the hash files are not being retained for a long enough duration, consider resizing the maximum size of the hash store by issuing the 'vserver cifs branchcache modify -hash-store-max-size' command.

#### Syslog Message

Size of the hashes in the hash store exceeded maximum size configured for Vserver %s.%s

#### Parameters

**vserver\_name** (STRING): Name of the Vserver where the problem is occurring.

**oldestKept\_timestamp** (STRING): Timestamp of the oldest kept hash file.

# hm events

## hm.alert events

### hm.alert.cleared

#### Severity

NOTICE

#### Description

This message occurs when a system health monitor alert is successfully cleared. Because health monitors evaluate the conditions dynamically, they might determine that the alerting condition is no longer valid. The alert is then removed from the system and this message is generated.

#### Corrective Action

(None).

#### Syslog Message

%s cleared by monitor %s

#### Parameters

**detailed\_info** (STRING): Specific information for the deleted alert.

**monitor** (STRING): Health monitor hm\_type value for this alert.

**alert\_id** (STRING): Alert ID for the alert.

**alerting\_resource** (STRING): Alert resource related to this alert.

### hm.alert.raised

#### Severity

ALERT

#### Description

This message occurs when a system health monitor alert is detected. Based on predefined policies, alerts are created dynamically as and when conditions are evaluated. The alert is then populated with the actual names of the resources causing it.

#### Corrective Action

Run the "system health alert show" command to see the alert that was created with this event. Follow the suggested actions listed in the "Corrective Actions" field of the alert.

#### Syslog Message

%s raised by monitor %s

#### Parameters

**detailed\_info** (STRING): Specific information for the created alert.

**monitor** (STRING): Health monitor hm\_type value for this alert.

**alert\_id** (STRING): Alert ID for the alert.

**alerting\_resource** (STRING): Alert resource related to this alert.

**probable\_cause\_description** (STRING): Description of what could have caused the alert.

**alert\_corrective\_action** (STRING): Instructions to correct the alerting condition.

**possible\_effect** (STRING): Possible effect on the system if the alert is not addressed.  
**additional\_info** (STRING): Additional information related to the alert.  
**tags** (STRING): Tags to categorize the alert.  
**acknowledge** (INT): True indicates the alert has been acknowledged.  
**acknowledger** (STRING): Text describing the user that acknowledged the alert.  
**suppress** (INT): True indicated the alert has been suppressed.  
**suppressor** (STRING): Text describing the user that suppressed the alert.  
**asup\_enable** (INT): True indicates that this alert will trigger an AutoSupport® message.

# host events

## host.app events

### host.app.component.error

#### Severity

ERROR

#### Description

This message occurs when a host encounters an error specific to an application component.

#### Corrective Action

For more information or assistance, contact NetApp technical support.

#### Syslog Message

Application component %s (UUID: %s) belonging to application %s (UUID: %s) on host %s running %s reported error %d: %s.

#### Parameters

**app\_name** (STRING): Name of the application.

**app\_uuid** (STRING): UUID of the application.

**component\_name** (STRING): Name of the application component.

**app\_component\_uuid** (STRING): UUID of the application component.

**host\_name** (STRING): Fully qualified domain name (FQDN) of the host.

**software\_product** (STRING): Name of the software product running on the host.

**host\_error\_num** (INT): Number of the error on the host.

**host\_error\_string** (STRING): Description of the error on the host.

### host.app.error

#### Severity

ERROR

#### Description

This message occurs when a host encounters an error specific to an application.

#### Corrective Action

For more information or assistance, contact NetApp technical support.

#### Syslog Message

Application %s (UUID: %s) on host %s running %s reported error %d: %s.

#### Parameters

**app\_name** (STRING): Name of the application.

**app\_uuid** (STRING): UUID of the application.

**host\_name** (STRING): Fully qualified domain name (FQDN) of the host.

**software\_product** (STRING): Name of the software product running on the host.

**host\_error\_num** (INT): Number of the error on the host.

**host\_error\_string** (STRING): Description of the error on the host.

# host.error events

## host.error

### Severity

ERROR

### Description

This message occurs when the host software encounters an error.

### Corrective Action

For more information or assistance, contact NetApp technical support.

### Syslog Message

Host %s running %s reported error %d: %s.

### Parameters

**host\_name** (STRING): Fully qualified domain name (FQDN) of the host.

**software\_product** (STRING): Name of the software product running on the host.

**host\_error\_num** (INT): Number of the error on the host.

**host\_error\_string** (STRING): Description of the error on the host.

# hotplug events

## hotplug.generic events

### hotplug.generic.busy\_not\_ready

#### Severity

ERROR

#### Description

This message occurs when a module replacement or removal command fails because the driver or system is busy. Hot plug functionality is not available at this time.

#### Corrective Action

Try the command again after a few minutes. If the system remains busy, it is an internal error. Contact NetApp technical support for assistance.

#### Syslog Message

%s command failed for module %s in slot %s of node %s because the driver or system is busy.

#### Parameters

**command** (STRING): Name of the command.

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert events

### hotplug.insert.badCard.type

#### Severity

ERROR

#### Description

This message occurs when the module type does not match the system's record for this slot.

#### Corrective Action

Use a replacement module of the same type.

#### Syslog Message

%s module insertion failed in slot %s of node %s because the module type does not match the system's record for this slot.

#### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert.badCard.ver

### Severity

ERROR

### Description

This message occurs when the module revision does not match the system's record for this slot.

### Corrective Action

Use a replacement module of the same board revision.

### Syslog Message

%s module insertion failed in slot %s of node %s because the module revision does not match the system's record for this slot.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert.illegal\_card

### Severity

NOTICE

### Description

This message occurs when the module inserted is not supported.

### Corrective Action

Use vendor-approved modules.

### Syslog Message

%s module insertion failed in slot %s of node %s because the module is not supported.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert.link\_not\_up

### Severity

ALERT

### Description

This message occurs when link on module fails to initialize after module is powered on.

### Corrective Action

Reinsert the module and try command again. If problem persists, contact NetApp technical support for assistance.



## Syslog Message

%s module insertion failed in slot %s of node %s because the link on module failed to initialize.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert.no\_power\_up

### Severity

ALERT

### Description

This message occurs when a module does not initialize after it is powered on.

### Corrective Action

Reinsert the module and try the command again. If the problem persists, contact NetApp technical support for assistance.

## Syslog Message

%s module insertion failed in slot %s of node %s because module failed to initialize.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.insert.success

### Severity

NOTICE

### Description

This message occurs when module insertion is successful.

### Corrective Action

(None).

## Syslog Message

%s module insertion is successful in slot %s of node %s.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.remove events

## hotplug.remove.success

### Severity

NOTICE

### Description

This message occurs when module removal is successful.

### Corrective Action

(None).

### Syslog Message

%s module removal is successful in slot %s of node %s.

### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

## hotplug.replace.events

### hotplug.replace.success

#### Severity

NOTICE

#### Description

This message occurs when module replacement is successful.

#### Corrective Action

(None).

#### Syslog Message

%s module replacement is successful in slot %s of node %s.

#### Parameters

**card** (STRING): Name of the module.

**slot** (STRING): Number of the slot.

**node** (STRING): Name of the node.

# http events

## http.cfg events

### http.cfg.csrf.max.tokens

#### Severity

NOTICE

#### Description

This message occurs when the maximum number of Cross-Site Request Forgery (CSRF) tokens has been reached. Web applications served from ONTAP®, such as ONTAP System Manager, use CSRF tokens to prevent unauthorized modifications to the cluster. When the token limit has been reached, new connections to these applications will not be able to make changes to the cluster configuration.

#### Corrective Action

You can fix the issue in a few different ways. You can disable CSRF protection by using the "system services web modify -csrf-protection-enabled false" command. You can increase the maximum number of tokens by using the "system services web modify -csrf-token-concurrent-limit <number>" command. You can wait until existing CSRF tokens expire. You can view and modify CSRF token settings by using the "system services web show" command.

#### Syslog Message

Maximum number of CSRF tokens (%d) has been reached.

#### Parameters

**limit** (INT): Maximum number of CSRF tokens that can exist at any given time.

# httpd events

## httpd.api events

### httpd.api.expired.session.limit

#### Severity

ERROR

#### Description

This message occurs when a Manage ONTAP® API request to an internal server expires due to configured session limits.

#### Corrective Action

Inspect the last error contained in the message to determine why the API request exceeded the configured session limits. The "security session show -interface ontapi" command may be used to display all active and pending API requests across the cluster. Run "security session limit show -interface ontapi" to verify the currently configured session limits.

#### Syslog Message

A Manage ONTAP® API request (%s) has expired after %lu seconds. Last error: %s

#### Parameters

**api\_name** (STRING): The name of the API.

**elapsed\_time** (LONGINT): The age of the API request (in seconds).

**last\_error** (STRING): Reason for the warning, indicating the limit category that was exceeded.

**purge\_time** (LONGINT): The purge time (in seconds) for outstanding API requests.

**trace\_tag** (LONGINTEX): The generated trace tag associated with the API request.

# hya events

## hya.dup\_vbn\_lookup events

### hya.dup\_vbn\_lookup

**Severity**

ERROR

**Description**

This event is generated if we detect redundant lookups of a PVBN without an intervening suspend.

**Corrective Action**

(None).

**Syslog Message**

%s %s %s %s %s %s %s %s %s %s %s %s

**Parameters**

- f2** (STRING): Third frame of backtrace.
- f3** (STRING): Fourth frame of backtrace.
- f4** (STRING): Fifth frame of backtrace.
- f5** (STRING): Sixth frame of backtrace.
- f6** (STRING): Seventh frame of backtrace.
- f7** (STRING): Eighth frame of backtrace.
- f8** (STRING): Ninth frame of backtrace.
- f9** (STRING): Tenth frame of backtrace.
- f10** (STRING): Eleventh frame of backtrace.
- f11** (STRING): Twelfth frame of backtrace.
- f12** (STRING): Thirteenth frame of backtrace.

### hya.dup\_vbn\_lookup.detail

**Severity**

ERROR

**Description**

This event is generated if we detect redundant lookups of a PVBN without an intervening suspend.

**Corrective Action**

(None).

**Syslog Message**

%s %s

**Parameters**

- description** (STRING): Description of failure.
- buffers** (STRING): Details of buffers that were predirtied.

# hybaggrcopy events

## hybaggrcopy.dstgeomismatch events

### hybAggrCopy.dstGeomMismatch

#### Severity

ERROR

#### Description

This message occurs when there is a mismatch between the disk geometries of source and destination hybrid aggregates.

#### Corrective Action

Satisfy the following conditions: Both aggregates are hybrid. The destination aggregate has at least the same number of HDD and SSD disks as the source aggregate. The disk size and type of each disk of the source aggregate matches those of the destination aggregate.

#### Syslog Message

Aggregate %s cannot be copied. Source and destination hybrid aggregate geometries do not match.

#### Parameters

**srcaggr** (STRING): Aggr copy source aggregate.

## hybaggrcopy.srcgeomismatch events

### hybAggrCopy.srcGeomMismatch

#### Severity

ERROR

#### Description

This message occurs when an aggregate cannot be copied. The issue is either that the source is a hybrid aggregate and the destination is not, or there is a mismatch between the disk geometries of the source and destination hybrid aggregates.

#### Corrective Action

Verify that the source and destination are both hybrid aggregates. Verify that the destination aggregate has at least the same number each of hard disk drives (HDDs) and solid state drives (SSDs) as the source aggregate. Verify that the disk size and type of each disk of the source aggregate matches those of the destination aggregate.

#### Syslog Message

Aggregate %s cannot be copied. Either the source is a hybrid aggregate but the destination is not, or the source and destination hybrid aggregate geometries do not match.

#### Parameters

**srcaggr** (STRING): Name of the source aggregate.

# ic events

## ic.hainterconnectdown events

### ic.HAInterconnectDown

#### Severity

ERROR

#### Description

This message occurs every hour when the high-availability (HA) interconnect is down.

#### Corrective Action

Base the corrective action on the number and type of HA interconnect links supported by the platform and the reason for which the interconnect is down. If links are down: Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabte the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. If links are disabled: Enable the links by using the "system ha interconnect link on" command. If a peer is not connected: Disable and reenabte the links one at a time by using the "system ha interconnect link off" and "system ha interconnect link on" commands. If the problem persists, contact NetApp technical support.

#### Syslog Message

HA interconnect: Interconnect down for %d minutes: %s

#### Parameters

**ic\_downtime** (INT): Amount of time (in minutes) the connection has been down.

**reason** (STRING): Reason that the HA interconnect is down, which might be due to links being down, links being disabled, or an unconnected peer.

## ic.hainterconnectlinkdown events

### ic.HAInterconnectLinkDown

#### Severity

ERROR

#### Description

This message occurs every hour when one of the high-availability (HA) interconnect links is down.

#### Corrective Action

Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabte the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. When one link is down on platforms with two available links, HA link redundancy is lost. However, HA services are still working and enabled on the secondary link.

## Syslog Message

HA interconnect: %s link #%d has been down for %d minutes.

## Parameters

**link\_type** (STRING): Type of link that is down. Possible values are external and backplane.

**link\_num** (INT): Number that identifies which link is down.

**link\_downtime** (INT): Amount of time (in minutes) the link has been down.

# ic.linkspeeddegraded events

## ic.linkSpeedDegraded

### Severity

ERROR

### Description

This message occurs when the high-availability (HA) interconnect detects a degraded link speed.

### Corrective Action

For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabling the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. If the problem persists, contact NetApp technical support.

## Syslog Message

HA interconnect: Link speed of link #%d degraded. Current: %s, expected: %s.

## Parameters

**link\_num** (INT): Number that identifies the link.

**link\_speed** (STRING): Current link speed.

**expected\_speed** (STRING): Link speed expected for the HA interconnect device.

# ic.linkstatuschange events

## ic.linkStatusChange

### Severity

INFORMATIONAL

### Description

This message occurs when the link status changes for a high-availability (HA) interconnect port.

### Corrective Action

Verify that both controllers in the HA pair are operational. For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabling the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. When one link is down on platforms with two available links, HA link redundancy is lost. However, HA services are still working and enabled on the secondary link.



## Syslog Message

HA interconnect: Port %s link is %s.

## Parameters

**port\_name** (STRING): Name of the port for which link status changed.

**link\_status** (STRING): Link status string: "up" or "down".

# ic.linkwidthdegraded events

## ic.linkWidthDegraded

### Severity

ERROR

### Description

This message occurs when the high-availability (HA) interconnect detects a degraded link width. This can affect the write performance of the system.

### Corrective Action

For externally connected links, make sure that the interconnect cables are connected properly and that small form-factor pluggable (SFP) modules, if applicable, are seated properly on both controllers. For internally connected links, disable and reenabling the links one at a time by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands. If the problem persists, contact NetApp technical support.

## Syslog Message

HA interconnect: Link width of link #%d degraded. Current: %dX, expected: %dX.

## Parameters

**link\_num** (INT): Number that identifies the link.

**link\_width** (INT): Current link width.

**expected\_width** (INT): Link width expected for the HA interconnect type.

# ic.unsupportcable events

## ic.unsupportCable

### Severity

ERROR

### Description

This message occurs when an unsupported cable is detected on the high-availability interconnect port. An unsupported cable might cause the interconnect to be unstable, resulting in unsynchronized logs.

### Corrective Action

Unplug the unsupported cable, and then plug in a supported cable. Search the support site "Hardware Universe" for information about supported high-availability cables for this platform.

## Syslog Message

Detected unsupported %s cable with part number %s.

## Parameters

**cable\_type** (STRING): Cable type.

**cable\_part** (STRING): Cable part number.

# ic.viif events

## ic.viif.initFailed

### Severity

ERROR

### Description

This message occurs when the driver fails to create a virtual interface.

### Corrective Action

Reboot the controller. Contact NetApp technical support if the problem persists.

### Syslog Message

HA interconnect: Virtual interface initialization failure.

### Parameters

(None).

# iclink events

## iclink.init events

### iclink.init.error

**Severity**

ALERT

**Description**

This message occurs when the HA interconnect fails to initialize.

**Corrective Action**

Reboot the system. If the problem persists, run diagnostics on the memory in the node.

**Syslog Message**

(None).

**Parameters**

**errorCode** (INT): The internal error code.

## iclink.probe events

### iclink.probe.error

**Severity**

ALERT

**Description**

This message occurs when the system cannot detect a High Availability (HA) interconnect card, making HA communication inoperable.

**Corrective Action**

Verify that the High Availability interconnect card is present and operational on the system. If this message persists, disable and reenable the link by using the (privilege: advanced) "system ha interconnect link off" and "system ha interconnect link on" commands.

**Syslog Message**

(None).

**Parameters**

**errorCode** (INT): Internal error code.

# icmp events

## icmp.ping events

### icmp.ping.throttle

#### Severity

NOTICE

#### Description

This message occurs when the system receives a burst of ICMP ping packet within a one-second interval. The packets are dropped by the system because these packets could be due to an ICMP attack.

#### Corrective Action

If the condition persists, determine the network source of the ICMP ping traffic and then block it by using a firewall rule. For assistance in determining the source, contact your network administrator.

#### Syslog Message

icmp\_input: The system dropped %d ping packets after limiting the packets to %d per second.

#### Parameters

**pingDropped** (INT): ICMP\_ECHO packets dropped due to throttling since the last time this message was generated.

**icmpLimit** (INT): Maximum number of ICMP packets per second.

# image events

## image.install events

### image.install.fail.aggregate

**Severity**

ERROR

**Description**

This message occurs when a software image installation fails, because one or more aggregates or volumes are not online.

**Corrective Action**

Ensure that aggregates and volumes are all online by using the "storage aggregate show -state !online" command and the "volume show -state !online -state !restricted" command.

**Syslog Message**

"Image installation fail: One or more aggregates are not online"

**Parameters**

(None).

### image.install.fail.mcluster

**Severity**

ERROR

**Description**

This message occurs when a software image installation fails because a MetroCluster(tm) switchover operation is in progress or already switched over.

**Corrective Action**

Use the "metrocluster show" command to verify that both the local and remote clusters are neither performing switchover nor in switchover mode.

**Syslog Message**

"MetroCluster switchover operation is in progress or already switched over."

**Parameters**

(None).

### image.install.fail.sfostate

**Severity**

ERROR

**Description**

This message occurs when a software image installation fails because the node is in takeover state.

**Corrective Action**

Ensure that all storage failover operations have been completed and that no nodes remain taken over by using the "storage failover show-giveback" command, and then try the installation again.

**Syslog Message**

"The node is taken over, cannot set to new image."

**Parameters**

(None).

**image.install.fail.sign****Severity**

ALERT

**Description**

This message occurs when the software installation process cannot verify the digital signature of the intended image.

**Corrective Action**

Replace the image, and then try the command again.

**Syslog Message**

"Image installation fail: Image digital signature check failed."

**Parameters**

(None).

# ipfw events

## ipfw.reachedmaxstates events

### ipfw.ReachedMaxStates

#### Severity

NOTICE

#### Description

This message occurs when the ipfw firewall fails to create a new dynamic state entry for a "keep-state" rule because the number of dynamic state entries has reached the maximum allowed value of 32768. The "keep-state" rule is used by the firewall to keep track of whether a connection is established. States are maintained by firewall for TCP, UDP, BGP, and ESP packets. This message occurs at most once every 60 seconds; it lists the most recent connections to reach the limit.

#### Corrective Action

(None).

#### Syslog Message

The ipfw firewall failed to create dynamic "keep-state" entry. Reason: %s, current # of entries: %d. Recent connections reaching this limit: %s

#### Parameters

**message** (STRING): Information about the failed ipfw dynamic state creation.

**current\_states\_held** (INT): Current total number of ipfw dynamic "keep state" entries held.

**cnn\_message** (STRING): Information about recent connections that reached the limit.

# ipmultipath events

## ipmultipath.init events

### ipmultipath.init.badChecksumType

**Severity**

ERROR

**Description**

This message occurs during IP Multipath session initialization, when an INIT packet with an unknown checksum type is received.

**Corrective Action**

The two end points are possibly running incompatible versions of Data ONTAP®. Run identical/compatible versions of Data ONTAP on the two end points and try again.

**Syslog Message**

Unknown checksum type (%d) in IP Multipath INIT packet.

**Parameters**

**chksum\_type** (INT): Checksum type received.

### ipmultipath.init.badVersion

**Severity**

ERROR

**Description**

This message occurs during IP Multipath session initialization, when an INIT packet with an unknown version number is received.

**Corrective Action**

The two end points are possibly running incompatible versions of Data ONTAP®. Run identical/compatible versions of Data ONTAP on the two end points and try again.

**Syslog Message**

Unknown version (%d) in IP Multipath INIT packet.

**Parameters**

**version** (INT): Version number received.

## ipmultipath.session events

### ipmultipath.session.noMemory

**Severity**

ERROR



## Description

This message occurs when an IP Multipath session is dropped due to failure to allocate memory.

## Corrective Action

The system is out of memory. Try to reduce load on the system. When memory becomes available, IP Multipath automatically reestablishes the session.

## Syslog Message

IP Multipath aborting session due to failure to allocate memory.

## Parameters

(None).

# ipmultipath.setup events

## ipmultipath.setup.bindFailed

### Severity

ERROR

### Description

This message occurs during IP Multipath connection setup, when the IP Multipath listener process cannot bind to any IP address of the source.

### Corrective Action

The source host information provided to IP Multipath is probably incorrect. Correct the configuration information and retry.

### Syslog Message

IP Multipath failed to bind to %s.

### Parameters

**src** (STRING): Host name on which the failure occurred.

## ipmultipath.setup.connFailed

### Severity

ERROR

### Description

This message occurs during IP Multipath connection setup, when it cannot connect to the remote end point.

### Corrective Action

The destination host information provided to IP Multipath is incorrect or the remote host is unreachable. Correct configuration information and retry. If the remote host is unreachable, IP Multipath automatically establishes the connection after the problem is corrected.

### Syslog Message

IP Multipath failed to connect to %s.

## Parameters

**dst** (STRING): Host name of the remote host.

# ipsec events

## ipsec.ib events

### ipsec.ib.bad.encrypt

#### Severity

ALERT

#### Description

This message occurs when an inbound IPsec packet cannot be successfully decrypted. Corrupt data is received.

#### Corrective Action

The IPsec SA database could be out of sync to cause this decryption failure. Use the "security ipsec policy modify -is-enabled" command to disable and then reenale the affected policy matching the source and destination IP address.

#### Syslog Message

IPsec: %s packet cannot be successfully decrypted.

#### Parameters

**sa2str** (STRING): IPsec Security Association (SA) in string format SA (Security Parameter Index, source IP address, destination IP address).

### ipsec.ib.no.lookup

#### Severity

ALERT

#### Description

This message occurs when an inbound IPsec packet cannot able to locate the IPsec Security Association (SA) from the SA database. The remote side is using an SA which does not exist in the local SA database.

#### Corrective Action

The IPsec SA database is out of sync. Use the "security ipsec policy modify -is-enabled" command to disable and then reenale the affected policy matching the source and destination IP address.

#### Syslog Message

IPsec: SA not found. SPI: %s Src IP: %s Dst IP: %s.

#### Parameters

**spi** (STRING): IPsec SA Security Parameter Index (SPI).

**srcAddr** (STRING): Source IP address.

**dstAddr** (STRING): Destination IP address.

### ipsec.ib.no.replay

**Severity**

NOTICE

**Description**

This message occurs when an inbound IPsec Security Association (SA) encounters a replay check failure. Number of out-of-order and retransmitted IPsec packet received. Possible causes include a denial-of-service (DoS) attack or traffic congestion.

**Corrective Action**

Possible traffic congestion or DoS attack.

**Syslog Message**

IPsec: %s replay packets detected. Possible DoS attack or traffic congestion.

**Parameters**

**sa2str** (STRING): IPsec SA in string format SA (Security Parameter Index, source IP address, destination IP address).

## ipsec.ib.policy.violation

**Severity**

ALERT

**Description**

This message occurs when an inbound IPsec packet is found to violate the security policy. It is either not IPsec protected or it is protected in a different way than the policy specifies.

**Corrective Action**

check the IPsec policy configuration on both ends to make sure that they match each other.

**Syslog Message**

IPsec: Inbound %s packet %s => %s security policy violation.

**Parameters**

**protocol** (STRING): Packet protocol.  
**srcAddr** (STRING): Source IP address.  
**dstAddr** (STRING): Destination IP address.

## ipsec.ob events

### ipsec.ob.no.acquire

**Severity**

NOTICE

**Description**

This message occurs when the kernel cannot acquire an IPsec Security Association (SA) to protect an outbound packet.

**Corrective Action**

Double check the IPsec policy configuration on both ends to make sure that they match each other.

**Syslog Message**

IPsec: Could not acquire SA for packet %s => %s.

**Parameters**

**srcAddr** (STRING): Source IP address.

**dstAddr** (STRING): Destination IP address.

# iscsi events

## iscsi.loginfailure events

### iscsi.loginFailure

#### Severity

NOTICE

#### Description

This event is issued when an initiator fails an iSCSI login attempt.

#### Corrective Action

Attempt to configure the iSCSI text keys on the initiator such that the iSCSI login succeeds. If the event occurs occasionally but does not affect stability or performance, no corrective action is needed but the event should be monitored. If the event persists and/or iSCSI application performance degrades, starting and stopping the iSCSI service and restarting any iSCSI applications may help clear the condition. Contact NetApp technical support.

#### Syslog Message

ISCSI: iSCSI login failure, '%s'

#### Parameters

**m** (STRING): A more detailed message describing the iSCSI login failure.

## iscsi.loginipnotauthorized events

### iscsi.loginIpNotAuthorized

#### Severity

NOTICE

#### Description

This message occurs when the system rejects an iSCSI login attempt because it originates from an unauthorized IP address.

#### Corrective Action

Run 'iscsi security show -address-ranges' to see if the initiator IP address has been added to the list of permitted IP addresses from where iSCSI logins are allowed.

#### Syslog Message

iSCSI login authorization failure. Initiator %s attempted to log in on Vserver %s, LIF %s from IP address %s

#### Parameters

**initiator\_name** (STRING): Name of the initiator that attempted to login.

**vserver** (STRING): Name of the vserver to which the unauthorized login attempt is made.

**lif** (STRING): Name of the LIF where the login was attempted.

**initiator\_ip\_address** (STRING): Source IP address of the initiator that attempted to login.

# iscsi.notice events

## iscsi.notice

### Severity

NOTICE

### Description

This event is issued when a loggable, non-error event occurs within the iSCSI subsystem.

### Corrective Action

(None).

### Syslog Message

ISCSI: %s

### Parameters

m (STRING): A more specific message describing the event.

# iscsi.session events

## iscsi.session.stateChanged

### Severity

NOTICE

### Description

This message occurs when iSCSI session status is changed.

### Corrective Action

If the identified iSCSI session is experiencing unexpected events (disconnecting, reconnecting, or connected), examine the link for excessive errors or latency.

### Syslog Message

iSCSI session state is changed to %s for the target %s (type: %s, address: %s). %s

### Parameters

**iscsi\_session\_state** (STRING): New iSCSI session state.

**iscsi\_target** (STRING): iSCSI target name.

**iscsi\_target\_type** (STRING): iSCSI target type.

**iscsi\_target\_address** (STRING): iSCSI target address.

**iscsi\_session\_state\_change\_reason** (STRING): Reason for the iSCSI session state change.

**iscsi\_timeout\_value** (INT): iSCSI timeout value.

# iscsi.warning events

## iscsi.warning

**Severity**

NOTICE

**Description**

This event is issued if the iSCSI subsystem detects an abnormal state.

**Corrective Action**

(None).

**Syslog Message**

ISCSI: %s

**Parameters**

**m** (STRING): A more specific message describing the state.



# isns events

## isns.enqueue events

### isns.enqueue.failed

#### Severity

ERROR

#### Description

This message occurs when the Internet Storage Name Service (iSNS) is not able to enqueue a new operation to the iSNS subsystem.

#### Corrective Action

The iSNS failed to enqueue an action into the work queue. This may be due to an overloaded system. Please consider moving the Vserver's Management LIF to another node for that Vserver in the cluster with a lighter work-load.

#### Syslog Message

The iSNS cannot enqueue an operation %s for Vserver %s.

#### Parameters

**event** (STRING): Event/Operation that was no enqueued.

**vserver** (STRING): Name of the Vserver

## isns.portal events

### isns.portal.info.unavailable

#### Severity

ERROR

#### Description

This message occurs when the Internet Storage Name Service (iSNS) is not able to retrieve the portal information required for iSNS registration.

#### Corrective Action

Ensure that at least one node hosting a IP data LIF for SAN is successfully loaded and operating correctly.

#### Syslog Message

The iSNS cannot retrieve the portal information for %s: %s.

#### Parameters

**vserver** (STRING): Name of the Vserver

**err\_msg** (STRING): Reported error.

## isns.reg events

## isns.reg.info.unavailable

### Severity

ERROR

### Description

This message occurs when the Internet Storage Name Service (iSNS) registration information is not valid or not available.

### Corrective Action

Check the Data ONTAP® iSNS configuration to ensure that all configuration information is present. Also check the cluster health for the impacted node to ensure that it is online.

### Syslog Message

The iSNS registration information for %s was not valid or not available: %s.

### Parameters

**vserver** (STRING): Name of the Vserver

**err\_msg** (STRING): Reported error.

## isns.reg.pdu.oversized

### Severity

ALERT

### Description

This message occurs when the the Internet Storage Name Service (iSNS) registration query packet data unit (PDU) length is larger than the supported maximum value. iSNS will not function properly until the problem is resolved.

### Corrective Action

Reconfigure the networking interfaces to reduce the portals until the iSNS registration query PDU is within its maximum value.

### Syslog Message

The iSNS registration query for %s failed due to the oversized iSNS query PDU length. Reduce the networking interface portals.

### Parameters

**vserver** (STRING): Name of the Vserver

## isns.reg.period.too\_small

### Severity

INFORMATIONAL

### Description

This message occurs when the Internet Storage Name Service (iSNS) requests a registration period that is too small. When this happens, the iSNS client service will be automatically disabled.

### **Corrective Action**

Reconfigure the server to increase the registration period.

### **Syslog Message**

The iSNS registration period for %s was too small (%d); increase the registration period.

### **Parameters**

**vserver** (STRING): Name of the Vserver

**period** (INT): Registration period received from the iSNS server

## **isns.socket.events**

### **isns.socket.conn.failed**

#### **Severity**

ERROR

#### **Description**

This message occurs when the TCP socket connection in the Internet Storage Name Service (iSNS) fails to connect to the iSNS server.

#### **Corrective Action**

The iSNS failed to create a socket for communicating with the iSNS server. Check the Data ONTAP® network configuration to ensure that the address for the iSNS server is in the routing group for all Vserver Management LIFs for the Vserver and that they are reachable via the clustered Data ONTAP "network ping" command.

#### **Syslog Message**

The iSNS TCP socket connection failed for %s to %s: %s

#### **Parameters**

**vserver** (STRING): Name of the Vserver

**ip\_address** (STRING): IP address for the iSNS server.

**err\_msg** (STRING): Reported network error.

### **isns.socket.setup.failed**

#### **Severity**

ERROR

#### **Description**

This message occurs when the Internet Storage Name Service (iSNS) fails to create a socket for communicating with the iSNS server.

#### **Corrective Action**

Check the Data ONTAP® network configuration to ensure that the address for the iSNS server is in the routing group for all Vserver Management LIFs for the Vserver and that they are reachable via the clustered Data ONTAP "network ping" command.

### **Syslog Message**

The iSNS for %s failed to create a socket to %s; check the Data ONTAP network configuration.

### **Parameters**

**vserver** (STRING): Name of the Vserver

**ip\_address** (STRING): IP address for the iSNS server.

# kern events

## kern.cron events

### kern.cron.dstError

**Severity**

ERROR

**Description**

This event is issued when the kernel cron daemon is unable to reschedule an internal process due to an error in handling of daylight savings time.

**Corrective Action**

(None).

**Syslog Message**

Daylight Saving Time change caused %s to be scheduled at an invalid time could not reschedule it, deleting further scheduling

**Parameters**

**owner** (STRING): The name of the process that was supposed to be scheduled

### kern.cron.reschedDst

**Severity**

NOTICE

**Description**

This event is issued when the kernel cron daemon reschedules an event to run at a newly scheduled time as a result of a daylight savings time change.

**Corrective Action**

(None).

**Syslog Message**

Daylight Savings Time change: rescheduling %s to run at %s

**Parameters**

**owner** (STRING): The name of the process that was scheduled

**time** (STRING): The text version of the time string at which the event was scheduled

## kern.log events

### kern.log.rotate

**Severity**

NOTICE

## Description

This event is issued when the system messages log is rotated to a new name. This event will be the first event in the new message file. When in takeover mode, this event is also issued for the partner node.

## Corrective Action

(None).

## Syslog Message

System %s (ID %s) is running %s

## Parameters

**hostname** (STRING): The host name of the appliance.

**serialno** (STRING): The serial number of the host.

**short\_version** (STRING): The short version string of the OnTap kernel, in the form 'NetApp Release NNNN'.

# kern.rc events

## kern.rc.errorMsg

### Severity

NOTICE

### Description

This event is issued when an error printf is issued by the kernel in the context of executing the `/etc/rc` file.

### Corrective Action

(None).

### Syslog Message

%s

### Parameters

**msg** (STRING): The text of the message

## kern.rc.msg

### Severity

NOTICE

### Description

This event is issued when a printf is issued by the kernel in the context of executing the `/etc/rc` file.

### Corrective Action

(None).

### Syslog Message

%s

## Parameters

**msg** (STRING): The text of the message

# kern.shutdown events

## kern.shutdown

### Severity

NOTICE

### Description

This event is issued to indicate a pending shutdown is underway.

### Corrective Action

(None).

### Syslog Message

System shut down because : "%s".

## Parameters

**type** (STRING): The type of shutdown (halt or reboot) that is underway.

## kern.shutdown.cantDeleteEntry

### Severity

ERROR

### Description

This event is issued when we detect that a halt or reboot is about to occur but a registry entry cannot be deleted

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**key** (STRING): The registry key

# kern.syslog events

## kern.syslog.drops

### Severity

ERROR

### Description

This event is generated when we detect that syslog messages were dropped before being logged. Syslog messages may be dropped if one or more Data ONTAP subsystems are generating many syslog messages

over short period of time. That can be an indicator that those subsystems are experiencing problems that need to be resolved.

### Corrective Action

Look in the syslog messages file at the messages prior to this message. If they all seem to be coming from a common subsystem, it may indicate that there is a problem with that subsystem. Attempt to resolve the problem with that subsystem.

### Syslog Message

Dropped %d syslog messages starting at %s due to lack of internal log buffers

### Parameters

**numDrops** (INT): The number of messages that have been dropped.

**timestamp** (STRING): A timestamp that indicates the time when the syslog messages started being dropped.

## kern.syslog.msg

### Severity

NOTICE

### Description

This event is generated when a kernel syslog message is detected.

### Corrective Action

(None).

### Syslog Message

%s

### Parameters

**msg** (STRING): The text of the message

**vfilerName** (STRING): The name of the vfiler that generated the message

## kern.syslogd events

### kern.syslogd.error

### Severity

ERROR

### Description

This message occurs when an event is generated to log errors from the syslog daemon.

### Corrective Action

(None).

### Syslog Message

syslog daemon error %s.



## Parameters

**ErrMsg** (STRING): Specific error generated by the syslog daemon.

# kern.timezone events

## kern.timezone.change

### Severity

NOTICE

### Description

This message occurs when the configured time zone changes.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**old\_tz** (STRING): Previous time zone.

**old\_tz\_offset** (INT): Seconds east of UTC of the previous time zone.

**new\_tz** (STRING): New time zone.

**new\_tz\_offset** (INT): Seconds east of UTC of the new time zone.

# kern.uptime events

## kern.uptime.filer

### Severity

NOTICE

### Description

This message occurs to display the uptime status of the controller. Typically, this message is generated once per hour.

### Corrective Action

(None).

### Syslog Message

%s

## Parameters

**msg** (STRING): Text message describing the current status of the controller.

**secs** (LONGINT): Number of seconds since the controller was booted.

**nfsOps** (LONGINT): Number of NFS operations that have been performed.

**cifsOps** (LONGINT): Number of CIFS operations that have been performed.

**httpOps** (LONGINT): Number of HTTP operations that have been performed.

**fcpOps** (LONGINT): Number of FCP operations that have been performed.

**iscsiOps** (LONGINT): Number of iSCSI operations that have been performed.

**nvmeFcOps** (LONGINT): Number of NVMe/FC operations that have been performed.

**nvmeTcpOps** (LONGINT): Number of NVMe/TCP operations that have been performed.

**nvmeRoceOps** (LONGINT): Number of NVMe/RDMA over Converged Ethernet(RoCE) operations that have been performed.

## kern.version events

### kern.version.change

#### Severity

NOTICE

#### Description

This message occurs when you reboot your system with a version of the Data ONTAP® kernel that is different from the one you ran previously.

#### Corrective Action

(None).

#### Syslog Message

Data ONTAP kernel version was changed from %s to %s.

#### Parameters

**oldver** (STRING): Previous version of the Data ONTAP kernel.

**newver** (STRING): Current version of the Data ONTAP kernel.

## kern.vm events

### kern.vm.mmap.return

#### Severity

NOTICE

#### Description

This message occurs when an attempt to create a new mapping in the virtual address space of a process has failed.

#### Corrective Action

(None).

#### Syslog Message

mmap(2) by %s (pid %d) for size %ld failed: %s, limit %ld, error %d.

#### Parameters

**process\_name** (STRING): Name of the process with the failed attempt.

**process\_id** (INT): ID of the process with the failed attempt.

**size** (LONGINT): Size in bytes of the memory map attempt.

**message** (STRING): Details about the failed memory map attempt.

**limit** (LONGINT): Size in bytes of the process's virtual memory limit.

**errno** (INT): The error code returned to the process.

# km events

## km.cluster events

### km.cluster.okm.hierarchy.bad

#### Severity

ALERT

#### Description

This message occurs in a MetroCluster(tm) configuration, when one or both clusters have been configured with the Onboard Key Manager, but one or both of the Onboard Key Manager key hierarchies are corrupted. Specifically, the 'hashed\_pass' object does not match the 'hashed\_pass\_!' object in at least one cluster.

#### Corrective Action

Reboot each of the nodes in the MetroCluster configuration.

#### Syslog Message

This cluster is part of a MetroCluster configuration. The Onboard Key Manager key hierarchy on one or both of the clusters is corrupted. Specifically, the 'hashed\_pass' object does not match the 'hashed\_pass\_!' object, part of the 'cluster\_kdb' table, in either cluster.

#### Parameters

(None).

## km.cmek events

### km.cmek.gckms.available

#### Severity

NOTICE

#### Description

This message occurs when a configured Google Cloud Key Management Service (GCKMS) that was previously reported as unavailable for key operations is now available.

#### Corrective Action

(None).

#### Syslog Message

The GCKMS with project ID: "%s", key ring location: "%s", key ring: "%s" and key name: "%s" configured is now available.

#### Parameters

**projectId** (STRING): Project ID.

**keyringLocation** (STRING): Location of the key ring.

**keyringName** (STRING): Name of the key ring.

**keyName** (STRING): Name of the key.

## km.cmek.gckms.keyVersionChg

### Severity

NOTICE

### Description

This message occurs when a change is made to the current key version associated with the key name of the configured Google Cloud Key Management Service (GCKMS). The key version might have been auto-rotated or manually changed at the GCKMS. Additionally, the top-level internal key protection key is successfully re-wrapped with the current key version of the GCKMS key name.

### Corrective Action

(None).

### Syslog Message

The version of the key "%s" owned by GCKMS project ID "%s", with key ring "%s" at "%s", has been changed from "%s" to "%s".

### Parameters

**keyName** (STRING): Name of the key.  
**projectId** (STRING): Project ID of the GCKMS.  
**keyringName** (STRING): Name of the key ring.  
**keyringLocation** (STRING): Location of the key ring.  
**oldKeyVersion** (STRING): Prior key version.  
**newKeyVersion** (STRING): New key version.

## km.cmek.gckms.notAvailable

### Severity

ALERT

### Description

This message occurs when a check for the availability of the configured Google Cloud Key Management Service (GCKMS) fails. The GCKMS might be down or there might be a network-related problem preventing communication. Without access to the GCKMS, the system might not be able to restore encryption keys needed to mount encrypted volumes. If the GCKMS is not available, then the next time this system boots, the failure to restore the keys might prevent the system from booting successfully or the encrypted volumes hosted on this system from coming online.

### Corrective Action

Verify that the GCKMS configuration is correct and verify the permissions and connectivity are correct. If the issue still persists, contact technical support.

### Syslog Message

The GCKMS with project ID: "%s", key ring location: "%s", key ring: "%s" and key name: "%s" is not available.

### Parameters

**projectId** (STRING): Project ID.  
**keyringLocation** (STRING): Location of the key ring.  
**keyringName** (STRING): Name of the key ring.  
**keyName** (STRING): Name of the key.

## km.cmek.init.failed

### Severity

ALERT

### Description

This message occurs when the system fails to initialize the key hierarchy required for volume encryption and that is protected by the Customer Managed Encryption Key (CMEK). Reasons for failure include failure to create the key hierarchy, incorrect CMEK configuration, identity, networking or permission issues in communicating with the CMEK hosted service and failure to store the CMEK wrapped key hierarchy in an external store like etcd.

### Corrective Action

Ensure that the configuration, identity, networking and permissions required to communicate with the CMEK hosted service are configured correctly. If the issue persists, contact technical support.

### Syslog Message

Initialization of the CMEK protected key hierarchy fails. Reason: "%s".

### Parameters

**reason** (STRING): Reason for failure.

## km.cmek.init.success

### Severity

NOTICE

### Description

This message occurs when the system successfully initializes and protects the key hierarchy with a Customer Managed Encryption Key (CMEK). The system is ready to support encrypted volumes.

### Corrective Action

(None).

### Syslog Message

CMEK protected key hierarchy initialization successfully completed.

### Parameters

(None).

## km.extkeysvr events

### km.extkeysvr.details

### Severity

NOTICE

### Description

This message occurs when a key is retrieved or restored from an external KMIP server. This is an informational message that describes the external KMIP server where the key is stored.

### Corrective Action

(None).

### Syslog Message

Key ID "%s" restored/retrieved from external key server "%s".

### Parameters

**keyId** (STRING): Key ID of the key being restored or retrieved.

**hostPort** (STRING): Host port details of the external KMIP server.

## km.keyretrieve events

### km.keyretrieve.failed

#### Severity

ALERT

#### Description

This message occurs when a node fails to retrieve a complete set of encryption keys from the configured key servers by the time the management process comes online. Should this node take ownership of an encrypted volume associated with a missing key before the key becomes available, the encrypted volume cannot come online.

#### Corrective Action

Use the "security key-manager external restore -node <node> -vserver <vserver>" command to manually restore the keys.

#### Syslog Message

Key auto-retrieve failed on node %s for Vserver %s (ID %u, UUID %s).

#### Parameters

**node** (STRING): Name of the node whose management process goes online.

**vserver** (STRING): Name of the Vserver that could not retrieve all of the keys.

**vserverId** (INT): Vserver ID of the Vserver.

**vserverUuid** (STRING): UUID of the related Vserver.

## km.keyserver events

### km.keyserver.available

#### Severity

NOTICE

#### Description

This message occurs when a check by the key manager for connectivity with a configured key management server indicates that the key management server that was previously unavailable is now available.

#### Corrective Action

(None).

## Syslog Message

The external key management server "%s" is now available for Vserver "%s".

## Parameters

**vserver** (STRING): Name of the Vserver where the key management server is configured.

**address** (STRING): IP address of the key management server.

## km.keyserver.notavailable

### Severity

ALERT

### Description

This message occurs when a check by the key manager for connectivity with a configured key management server indicates that the key management server is not available. The key management server might be down, there might be a network-related problem preventing communication with the key server, or the security certificates used to authenticate with the key management server might have expired. Without access to the key management server, the node might not be able to restore authentication keys needed to unlock NSE drives or encryption keys needed to mount encrypted volumes. If the key management server is not available the next time this node boots, then the failure to restore the keys might prevent the node from booting successfully or prevent the encrypted volumes hosted on this node from coming online.

### Corrective Action

Ensure that the key management server is online using the server's management interface. If the server is online, use the "network ping -lif <lif> -vserver <vserver> -destination <address>" command to verify that the node can communicate with the server. If this command indicates that the server is available, use the "security key-manager external show -vserver <vserver>" command to get the certificate names, and then check whether they are expired with the "security certificate show-user-installed -vserver <vserver> -cert -name <cert>" command. If so, install the updated certificates with the "security certificate install" and "security key-manager external modify" commands.

## Syslog Message

The external key management server "%s" is not available for Vserver "%s", status: "%s".

## Parameters

**vserver** (STRING): Name of the Vserver where the key management server is configured.

**address** (STRING): IP address of the key management server.

**keyserver\_status** (STRING): Status of the key management server.

## km.mcc events

### km.mcc.okm.hierarchy.bad

### Severity

ALERT

### Description

This message occurs in a MetroCluster(tm) configuration, when one or both clusters have been configured with the Onboard Key Manager, but one or both of the Onboard Key Manager key hierarchies are corrupted. Specifically, the 'hashed\_pass' object does not contain the same value in each Onboard Key Manager key hierarchy.

**Corrective Action**

Reboot each of the nodes in the MetroCluster configuration.

**Syslog Message**

This cluster is part of a MetroCluster configuration. The Onboard Key Manager key hierarchy on one or both of the clusters is corrupted. The replication subsystem has failed to update the 'hashed\_pass' field in the 'cluster\_kdb' table.

**Parameters**

(None).

**km.mcc.okmkey.mismatch****Severity**

ALERT

**Description**

This message occurs in a NetApp MetroCluster configuration, where both peer clusters are configured with Onboard Key Manager (OKM) and the lists of OKM hierarchy keys on the peer clusters are not consistent. A mismatch in the OKM hierarchy keys can lead to a failure during a reboot or future switchover/switchback event.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

This cluster is part of a MetroCluster configuration. The list of OKM hierarchy keys on one or both of the clusters is potentially corrupted.

**Parameters**

(None).

**km.mcc.svmkek.missing.local****Severity**

ALERT

**Description**

This message occurs in a NetApp MetroCluster configuration, where both peer clusters are configured with Onboard Key Manager (OKM) and an SVM key encryption key is created on the remote peer, but due to various conditions, is not replicated to the local cluster. Creation of a new SVM key encryption key (KEK) on the local cluster is prevented to avoid failure during a reboot or future switchover/switchback event.

**Corrective Action**

Contact Contact NetApp technical support. for assistance.

**Syslog Message**

This cluster is part of a MetroCluster configuration. An SVM key encryption key created on the remote cluster has not been replicated to the local cluster.



**Parameters**

(None).

## km.okmdb events

**km.okmdb.read.failed****Severity**

ALERT

**Description**

This message occurs when a read of the Onboard Key Manager's database file fails. This may occur when the Onboard Key Manager has been configured to use a USB based mount-point and the USB device has not been mounted on the node.

**Corrective Action**

Check that the USB device is attached to the node.

**Syslog Message**

The system is not able to read from the Onboard Key Manager's database file.

**Parameters**

(None).

**km.okmdb.write.failed****Severity**

ALERT

**Description**

This message occurs when a write to the Onboard Key Manager's database file fails. This may occur when the Onboard Key Manager has been configured to use a USB based mount-point and the USB device has either not been mounted on the node or the USB device is write-protected.

**Corrective Action**

Check that the USB device is attached to the node and that it is not write-protected.

**Syslog Message**

The system is not able to write to the Onboard Key Manager's database file.

**Parameters**

(None).

## km.onboard events

**km.onboard.ccmode.wrongpass****Severity**

NOTICE

## Description

This message occurs when a cluster administrator supplies the wrong Onboard Key Manager cluster-wide passphrase more than three times in a row while attempting to enable Common Criteria (CC) mode.

## Corrective Action

(None).

## Syslog Message

The Onboard Key Manager cluster-wide passphrase has been entered incorrectly %s times by user "%s" while attempting to enable Common Criteria mode.

## Parameters

**username** (STRING): Name of the user who issued the command to enable CC mode.

**failure\_count** (STRING): Number of sequential times that the cluster-wide passphrase has been entered incorrectly.

# km.restore events

## km.restore.invalid.key

### Severity

ALERT

### Description

This message occurs when the key manager attempts to restore a key which has an invalid length.

### Corrective Action

Run the "security key-manager restore" command to restore keys from external key servers.

### Syslog Message

ONTAP has attempted to restore an %s key which has an invalid key length. Vserver: %s, Key ID: %s, Key server: %s.

### Parameters

**keyType** (STRING): Key type (NSE, AES, XTS) that has invalid length.

**vserverName** (STRING): Vserver name.

**keyId** (STRING): Key ID of the key being restored.

**keyServer** (STRING): Key server from which the key is being restored.

# km.run events

## km.run.external.setup

### Severity

ERROR

### Description

This message occurs when a node successfully performs a join operation to join a cluster, but the other nodes in the cluster have the External Key Manager configured and the new node does not. The new node cannot perform External Key Manager-related tasks, such as creating encrypted volumes, until it has the

External Key Manager configured.

### Corrective Action

Set up the External Key Manager on the new node by using the "security key-manager setup" command.

### Syslog Message

External Key Manager setup required after the cluster "%s" operation for node "%s", serial "%s".

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**node** (STRING): Name of the node that successfully joined the cluster.

**serial** (STRING): System serial number of the node that successfully joined the cluster.

## km.run.extrnl.enable.needed

### Severity

ALERT

### Description

This message occurs when a cluster in a MetroCluster(tm) configuration has configured external key management via the "security key-manager external enable -key-servers <ip\_address:port> -client-cert <client\_cert\_name> -server-ca-certs <server\_ca\_cert\_name>" command. This is an alert that the external key management configuration on the two clusters is not consistent, and that the "security key-manager external enable -key-servers <ip\_address:port> -client-cert <client\_cert\_name> -server-ca-certs <server\_ca\_cert\_name>" command" should be run on the local cluster or peer cluster, specifying the same set of key servers. Otherwise, encrypted volumes and NSE drives will not come online after a future switchover or switchback event. This message will be generated on both the local cluster and the peer cluster.

### Corrective Action

Run the "security key-manager external enable" command on the local cluster or the peer cluster (identified in the "cluster" parameter), specifying the same set of configured key servers, to make the external key management consistent.

### Syslog Message

The external key management configuration is not consistent between the local cluster and the peer cluster of the MetroCluster configuration. This can lead to a failure during a future switchover or switchback event. Synchronize the configuration on the %s cluster.

### Parameters

**cluster** (STRING): Local cluster or peer cluster on which to run the command.

## km.run.extrnl.setup.needed

### Severity

ALERT

### Description

This message occurs when a cluster in a MetroCluster(tm) configuration has configured external key management via the "security key-manager setup" command. This is an alert that the external key management configuration on the two clusters is not consistent, and that the "security key-manager setup" command should be run on the local cluster or peer cluster. Otherwise, encrypted volumes and NSE drives

will not come online after a future switchover or switchback event. This message will be generated on both the local cluster and the peer cluster.

### Corrective Action

Run the "security key-manager setup" command on the local cluster or the peer cluster (identified in the "cluster" parameter) to make the external key management configurations consistent.

### Syslog Message

The external key management configuration is not consistent between the local cluster and the peer cluster of the MetroCluster configuration. This can lead to a failure during a future switchover or switchback event. Synchronize the configuration on the %s cluster.

### Parameters

**cluster** (STRING): Local cluster or peer cluster on which to run the command.

## km.run.onboard.setup

### Severity

ERROR

### Description

This message occurs when a node successfully performs a join operation to join a cluster, but the other nodes in the cluster have the Onboard Key Manager configured and the new node does not. The new node cannot perform the Onboard Key Manager-related tasks, such as creating encrypted volumes, until it has the Onboard Key Manager configured.

### Corrective Action

Set up the Onboard Key Manager on the new node by using the "security key-manager setup" command.

### Syslog Message

The Onboard Key Manager setup required after the cluster "%s" operation for node "%s", serial "%s".

### Parameters

**operation\_type** (STRING): Type of cluster operation.

**node** (STRING): Name of the node that successfully joined the cluster.

**serial** (STRING): System serial number of the node that successfully joined the cluster.

## km.run.onboard.setup.needed

### Severity

ALERT

### Description

This message occurs when a cluster in a MetroCluster(tm) configuration has either configured the Onboard Key Manager via the "security key-manager setup" command or has reconfigured the cluster passphrase via the "security key-manager update-passphrase" command. This is an alert that the onboard key hierarchies on the two clusters are not consistent, and that the "security key-manager setup -sync -metrocluster-config yes" command should be run on the local cluster or peer cluster. Otherwise, a future switchover or switchback event could fail. This message will be generated on both the local cluster and the peer cluster.

## Corrective Action

Run the "security key-manager setup -sync-metrocluster-config yes" command on the local cluster or the peer cluster (identified in the "cluster" parameter) to make the two onboard key hierarchies consistent.

## Syslog Message

The Onboard Key Manager configuration is not consistent between the local cluster and the peer cluster of the MetroCluster configuration. This can lead to a failure during a future switchover or switchback event. Synchronize the configuration on the %s cluster.

## Parameters

**cluster** (STRING): Local cluster or peer cluster on which to run the command.

## km.run.onboard.sync.needed

### Severity

ALERT

### Description

This message occurs when a cluster in a MetroCluster(tm) configuration has either configured the Onboard Key Manager via the "security key-manager onboard enable" command or has reconfigured the cluster passphrase via the "security key-manager onboard update-passphrase" command. This is an alert that the onboard key hierarchies on the two clusters are not consistent, and that the "security key-manager onboard sync" command should be run on the local cluster or peer cluster. Otherwise, a future switchover or switchback event could fail. This message will be generated on both the local cluster and the peer cluster.

## Corrective Action

Run the "security key-manager onboard sync" command on the local cluster or the peer cluster (identified in the "cluster" parameter) to make the two onboard key hierarchies consistent.

## Syslog Message

The Onboard Key Manager configuration is not consistent between the local cluster and the peer cluster of the MetroCluster configuration. This can lead to a failure during a future switchover or switchback event. Synchronize the configuration on the %s cluster.

## Parameters

**cluster** (STRING): Local cluster or peer cluster on which to run the command.

## km.volume events

### km.volume.mount.fail

### Severity

ALERT

### Description

This message occurs during the recovery of the key manager when a volume is placed online, but fails to mount. Volumes need to be mounted manually.

## Corrective Action

The volume needs to be mounted manually with the provided junction using the "volume mount" command.

## Syslog Message

Vserver "%s" failed to mount volume: "%s" to junction path: "%s" during key manager recovery.

## Parameters

**vserver** (STRING): Name of the Vserver that owns the volumes being placed online.

**volume** (STRING): Name of the the volume being placed online.

**junction** (STRING): Name of the junction path to mount.

## km.volume.mount.mismatch

### Severity

ALERT

### Description

This message occurs during recovery of the key manager when a junction path cannot identify the volume to mount. Pair the supplied junction with its volume and mount manually.

### Corrective Action

The volumes are online but need to be mounted manually using the "volume mount" command.

## Syslog Message

During health monitor recovery of the key manager on Vserver "%s", the junction path "%s" cannot be associated with the volume to mount.

## Parameters

**vserver** (STRING): Name of the Vserver that owns the volume that could not be mounted.

**junction** (STRING): Name of the junction path to mount.

# kmip2 events

## kmip2.cannot events

### kmip2.cannot.import.NSE.key

#### Severity

ALERT

#### Description

This message occurs when the KMIP2 client cannot import an NSE key from a KMIP key server. A key was retrieved from a KMIP key server, but an error occurred when trying to add the key to an ONTAP® data structure.

#### Corrective Action

Run the "security key-manager restore" command to restore keys from external key servers.

#### Syslog Message

KMIP2 client failed to import NSE key from KMIP key server, error: %d

#### Parameters

**err** (INT): Error code from system call to sysctlbyname().

### kmip2.cannot.import.NVE.key

#### Severity

ALERT

#### Description

This message occurs when the KMIP2 client cannot import an NVE key from a KMIP key server. A key was retrieved from a KMIP key server, but an error occurred when trying to add the key to an ONTAP® data structure.

#### Corrective Action

Run the "security key-manager restore" command to restore keys from external key servers.

#### Syslog Message

KMIP2 client failed to import NVE key from KMIP key server, error: %d

#### Parameters

**err** (INT): Error code from system call to sysctlbyname().

### kmip2.cannot.read.client.cert

#### Severity

ALERT

#### Description

This message occurs when the client certificate cannot be read. On a Cloud ONTAP® system, this message indicates one of the following: 1) A request to create an encrypted aggregate will not succeed

because the key cannot be sent to the external KMIP server, or 2) The request to retrieve a key from the external KMIP server will fail. In this case, the encrypted aggregate will not be able to be mounted. On a NSE system, this message indicates one of the following: 1) A request to re-key a NSE drive will not succeed because the key cannot be sent to the external KMIP server, or 2) The ability to read/write data to/from a NSE drive might be impacted if power to the drive or to the disk shelf is lost.

### **Corrective Action**

Examine the KMIP client certificate and ensure that the certificate has been created properly.

### **Syslog Message**

KMIP Failed to read client certificate with error: %s

### **Parameters**

**err** (STRING): KMIP error code.

## **kmip2.cannot.read.private.key**

### **Severity**

ALERT

### **Description**

This message occurs when the client private key certificate can't be read. On a Cloud ONTAP® system, this message indicates one of the following: 1) A request to create an encrypted aggregate will not succeed because the key cannot be sent to the external KMIP server, or 2) The request to retrieve a key from the external KMIP server will fail. In this case, the encrypted aggregate will not be able to be mounted. On a NSE system, this message indicates one of the following: 1) A request to re-key a NSE drive will not succeed because the key cannot be sent to the external KMIP server, or 2) The ability to read/write data to/from a NSE drive might be impacted if power to the drive or to the disk shelf is lost.

### **Corrective Action**

Examine the KMIP client private certificate and ensure that the certificate has been created properly.

### **Syslog Message**

KMIP Failed to read client private certificate with error: %s

### **Parameters**

**err** (STRING): KMIP error code.

## **kmip2.invalid events**

### **kmip2.invalid.key.length**

### **Severity**

ALERT

### **Description**

This message occurs when ONTAP cannot push a key to a KMIP key server because the key has an invalid length.

### **Corrective Action**

(None).



## Syslog Message

ONTAP failed to push an %s key to a KMIP key server due to invalid key length. Vserver Name: %s, Key ID: %s.

## Parameters

**keyType** (STRING): Key type (NSE, AES, XTS) that has invalid length.

**vserverName** (STRING): Vserver name.

**keyID** (STRING): Key ID of key being pushed to key server.

# kmip2.ssl events

## kmip2.ssl.cannot.connect

### Severity

ALERT

### Description

This message occurs when a KMIP SSL/TLS connection error is detected in either a Cloud ONTAP system with encrypted aggregates or a NSE system. On a Cloud ONTAP® system, this message indicates one of the following: 1) A request to create an encrypted aggregate will not succeed because the key cannot be sent to the external KMIP server, or 2) The request to retrieve a key from the external KMIP server will fail. In this case, the encrypted aggregate will not be able to be mounted. On a NSE system, this message indicates one of the following: 1) A request to re-key a NSE drive will not succeed because the key cannot be sent to the external KMIP server, or 2) The ability to read/write data to/from a NSE drive might be impacted if power to the drive or to the disk shelf is lost.

### Corrective Action

Examine the KMIP certificates and ensure that the certificates have been created properly.

### Syslog Message

Unable to make SSL/TLS connection to KMIP server. Error: %s

### Parameters

**err** (STRING): KMIP error code with details.

# krb events

## krb.config events

### krb.config.updateFail

**Severity**

ERROR

**Description**

The CIFS server's security settings could not be written to the configuration file on this node.

**Corrective Action**

Modify any of the 'vserver cifs security' fields using the 'vserver cifs security modify' command, which will attempt to write to the configuration file again on every node. Change the setting back to its original value if desired. Afterward, verify that the event log does not contain any new instances of this error.

**Syslog Message**

The CIFS server's security settings could not be written to the configuration file on this node.

**Parameters**

(None).

# ktls events

## ktls.badauth events

### ktls.badAuth

#### Severity

NOTICE

#### Description

This message occurs when the system receives a TLS message with bad 'authentication' data. The connection is closed. It might indicate that this system is under attack.

#### Corrective Action

ONTAP software automatically closes the connection from which this incorrect authentication information was received. If this message is repeatedly generated from the same internet address, it might indicate that this system is under attack.

#### Syslog Message

ONTAP received a TLS message with incorrect authentication at local address %s local port %d from remote address %s remote port %d via IPspace %d.

#### Parameters

**localAddr** (STRING): Full internet address of the local end of this session.

**localPort** (INT): Local port number that was being reached.

**remoteAddr** (STRING): Full internet address of the remote end of this session.

**remotePort** (INT): Remote port number that sent the erroneous data.

**IPspace** (INT): IPspace in which the remote address was reached.

## ktls.cnxnhandshakelimit events

### ktls.cnxnHandshakeLimit

#### Severity

NOTICE

#### Description

This message occurs when the concurrent TLS connection handshake limit is reached. Any new TLS connection request will fail and an error status will be returned.

#### Corrective Action

Reduce the encrypted connection load on ONTAP to avoid reaching the maximum concurrent TLS handshake limit.

#### Syslog Message

ONTAP reached the maximum limit of %d concurrent TLS connection handshakes. If this limit is reached, subsequent TLS connections will fail.

## Parameters

**maxTLSConnHandles** (INT): Maximum number of concurrent TLS connection handshakes.

# ktls.failed events

## ktls.failed

### Severity

NOTICE

### Description

This message occurs when several successive Transport Layer Security (TLS) connections with a remote system have failed.

### Corrective Action

Corrective actions will necessarily depend on what condition is being reported.

### Syslog Message

"The TLS connections have failed several times with remote host '%s' in IPspace '%s', for which the latest reason given is: %s."

### Parameters

**peer** (STRING): Network address of the remote peer.

**ipsname** (STRING): IPspace within which the network address is defined.

**reason** (STRING): Reason for failure.

# ktls.health events

## ktls.health.actionQGrowing

### Severity

NOTICE

### Description

This message occurs when the depth of the action queue in the ktlsd subsystem is increasing in the majority of time intervals. It is a potential indication of back-pressure accumulating within ktlsd or the kernel.

### Corrective Action

(None).

### Syslog Message

"The depth of the action queue in the ktlsd subsystem has trended upward over the last "%d" intervals of "%d" seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.actionQHigh

### Severity

NOTICE

### Description

This message occurs when the depth of the action queue in the ktlsd subsystem accounts for a significant percentage of our total available handshake slots. It is a potential indication of back-pressure accumulating within ktlsd or the kernel.

### Corrective Action

(None).

### Syslog Message

"The depth of the action queue in the ktlsd subsystem has been at an elevated level over the last %d intervals of %d seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.backoffGrowing

### Severity

NOTICE

### Description

This message occurs when the number of backoffs in the ktlsd subsystem is increasing in the majority of time intervals. It is most likely an indication that handshake timeouts are occurring frequently.

### Corrective Action

(None).

### Syslog Message

"The number of self-imposed backoffs in the ktlsd subsystem has trended upward over the last "%d" intervals of "%d" seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.highStress

### Severity

ALERT

### Description

This message occurs when a combination of indicators suggest high stress within the Kernel Transport Layer Security (KTLS) subsystem. This subsystem does not perform adequately as a consequence.

## Corrective Action

There might be a variety of potential corrective actions. The first course of action should be to check for other `ktls.health.*` events around the same time. These might have their own guidance. In general it is likely that the resolution involves reducing the inbound load on KTLS, which requires different steps depending on the feature.

## Syslog Message

"The KTLS subsystem is exhibiting signs of high stress."

## Parameters

(None).

## `ktls.health.inactiveGrowing`

### Severity

NOTICE

### Description

This message occurs when the total length of time handshakes spend in an active state has been increasing in the majority of time intervals.

### Corrective Action

(None).

### Syslog Message

"The inactive period for TLS handshakes has trended upward over the last "%d" intervals of "%d" seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## `ktls.health.latencyGrowing`

### Severity

NOTICE

### Description

This message occurs when the latency of handshakes has been increasing in the majority of time intervals.

### Corrective Action

(None).

### Syslog Message

"The latency of TLS handshakes has trended upward over the last "%d" intervals of "%d" seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.lowCompletion

### Severity

ERROR

### Description

This message occurs when a significant percentage of handshakes initiated in the ktlsd subsystem over an interval of time do not succeed.

### Corrective Action

There might be a variety of potential corrective actions. Any ktls.failed events generated around the same time contain a "reason" parameter which provides a good starting point for probable causes of this condition.

### Syslog Message

"A high percentage of handshakes initiated in the ktlsd subsystem over the last "%d" intervals of "%d" seconds have not completed successfully."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.startLatencyHigh

### Severity

NOTICE

### Description

This message occurs when the amount of time handshakes are spending in the start queue in the ktlsd subsystem is at an elevated level over some time intervals. It is a potential indication of back-pressure accumulating within ktlsd.

### Corrective Action

(None).

### Syslog Message

"The amount of time handshakes have spent in the start queue of the ktlsd subsystem has been at an elevated level over the last "%d" intervals of "%d" seconds."

### Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.startQGrowing

### Severity

NOTICE

### Description

This message occurs when the depth of the start queue in the ktlsd subsystem is increasing in the majority of time intervals. It is a potential indication of back-pressure accumulating within ktlsd.

## Corrective Action

(None).

## Syslog Message

"The depth of the start queue in the ktlsd subsystem has trended upward over the last "%d" intervals of "%d" seconds."

## Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.startQHigh

### Severity

NOTICE

### Description

This message occurs when the depth of the start queue in the ktlsd subsystem accounts for a significant percentage of total available handshake slots. It is a potential indication of back-pressure accumulating within ktlsd.

## Corrective Action

(None).

## Syslog Message

"The depth of the start queue in the ktlsd subsystem has been at an elevated level over the last "%d" intervals of "%d" seconds."

## Parameters

**intervalQuantity** (INT): Number of intervals over which the problem was observed.

**intervalLength** (INT): Length of the observation intervals for this problem, in seconds.

## ktls.health.stress

### Severity

ERROR

### Description

This message occurs when a combination of indicators suggest moderate stress within the Kernel Transport Layer Security (KTLS) subsystem.

## Corrective Action

There might be a variety of potential corrective actions. The first course of action should be to check for other ktls.health.\* events around the same time. These might have their own guidance. In general it is likely that the resolution involves reducing the inbound load on KTLS, which requires different steps depending on the feature.

## Syslog Message

"The KTLS subsystem is exhibiting signs of moderate stress."



**Parameters**

(None).

# largevol events

## largevol.remount events

### largevol.remount.nfsclient

#### Severity

NOTICE

#### Description

This message occurs when the "is-large-size-enabled" property is set on a FlexVol volume or Flexgroup volume, which is NFS mounted for client access. Setting this field enables volumes to support large files, from 16TB to 128TB. If the volume has already been mounted via NFS, then the NFS client caches the older maximum file size, which is 16TB. For support of files up to 128TB to take effect, the NFS client needs to remount the volume. This issue does not occur if automount has been configured on NFS client to mount the volumes.

#### Corrective Action

To support 128TB as the maximum file size on the NFS client, stop the applications using the NFS mount, remount the FlexVol volume in the NFS client and then restart the applications. If the issue still persists, Contact NetApp technical support.

#### Syslog Message

Remount the volume on the NFS client.

#### Parameters

**vvolname** (STRING): Volume name.

**vserver** (STRING): Vserver name.

# Idap events

## Idap.false events

### Idap.false.configs.removed

**Severity**

NOTICE

**Description**

This message occurs when the last system to be upgraded in a cluster removes LDAP configurations that were marked as disabled.

**Corrective Action**

(None).

**Syslog Message**

All LDAP configurations with the "-client-enabled" parameter set to "false" were removed.

**Parameters**

(None).

# lib events

## lib.nvtrace events

### lib.nvtrace.syslog.relog.alert

#### Deprecated

This event is deprecated as of ONTAP 9.1.

#### Severity

ALERT

#### Description

This message occurs when the system sorts the syslog messages when syslog is not yet initialized, and the panic messages are relogged.

#### Corrective Action

(None).

#### Syslog Message

Relog syslog %s.

#### Parameters

**data** (STRING): Syslog data in NVRAM(Non Volatile RAM).

# lic events

## lic.cappool events

### lic.cappool.reassign.abort

#### Severity

NOTICE

#### Description

This message occurs when the system aborts an operation of reassigning nodes to a different capacity pool license, because another reassignment operation to a different capacity pool license in the same request has failed.

#### Corrective Action

Look for the "lic.cappool.failed" event that occurred about the same time, which caused the reassignment operation to abort.

#### Syslog Message

"Capacity pool license reassignment is aborted for nodes ("%s") to serial number "%s" due to failure to reassign other nodes to serial number "%s"."

#### Parameters

**aborted\_node\_names** (STRING): Aborted list of node names.

**aborted\_pool\_serial** (STRING): Aborted capacity pool license serial number.

**failed\_pool\_serial** (STRING): Another reassignment operation to this capacity pool serial number has failed.

### lic.cappool.reassign.done

#### Severity

NOTICE

#### Description

This message occurs when the system successfully completes an operation of reassigning nodes to a different capacity pool license.

#### Corrective Action

(None).

#### Syslog Message

"Capacity pool license reassignment successfully completed for nodes ("%s") to license serial number "%s"."

#### Parameters

**node\_names** (STRING): List of node names that are reassigned to the new capacity pool license.

**new\_pool\_serial** (STRING): Capacity pool license serial number to which the nodes are reassigned.

## lic.cappool.reassign.done.info

### Severity

NOTICE

### Description

This message occurs when the system successfully completes an operation of reassigning nodes to a different capacity pool license with with a non-critical error, which does not affect the outcome of the ultimate operation.

### Corrective Action

(None).

### Syslog Message

"Capacity pool license reassignment successfully completed for nodes ("%s") to license serial number "%s" with non-critical error "%s"."

### Parameters

**node\_names** (STRING): List of node names that are reassigned to the new capacity pool license.

**new\_pool\_serial** (STRING): Capacity pool license serial number to which the nodes are reassigned.

**non\_critical\_error** (STRING): Information about the non-critical error encountered during the reassignment operation.

## lic.cappool.reassign.failed

### Severity

ERROR

### Description

This message occurs when operation of reassigning nodes to a different capacity pool license fails.

### Corrective Action

Resolve the error, which should clearly indicate why the operation has failed. Follow the instructions shown on the ONTAP Select Deploy utility from which the operation was initiated.

### Syslog Message

"Capacity pool license reassignemnt to license serial number "%s" failed for nodes ("%s") with error: %s"

### Parameters

**failed\_pool\_serial** (STRING): Capacity pool license serial number to which the nodes failed to be reassigned.

**failed\_node\_names** (STRING): List of node names that failed to be reassigned to the new capacity pool license.

**error** (STRING): Reason for failure.

## lic.cappool.reassign.start

### Severity

NOTICE

## Description

This message occurs when the system starts an operation of reassigning nodes to a different capacity pool license.

## Corrective Action

(None).

## Syslog Message

"Capacity pool license reassignment started for nodes ("%s") to license serial number "%s"."

## Parameters

**node\_names** (STRING): List of node names that are reassigned to the new capacity pool license.

**new\_pool\_serial** (STRING): Capacity pool license serial number to which the nodes are reassigned.

# lic.lm events

## lic.lm.server.not.configured

### Severity

ALERT

### Description

This message occurs when ONTAP® software is running in the capacity pools licensing model and there is no License Manager configuration available. Without the License Manager configuration, the system cannot create or expand an aggregate.

### Corrective Action

Find the IP address of the ONTAP Select Deploy instance from which this ONTAP system was created. Use that IP address to manually configure the License Manager by using the "system license license-manager modify -host ip\_address\_of\_license\_manager" command. Verify that the correct License Manager is configured by using the "system license license-manger show" command. Use the "system license license-manager check" command to determine whether the License Manager is accessible. These commands are accessible through diagnostic privileges.

### Syslog Message

"License Manager has not been configured for capacity pools."

### Parameters

(None).

# license events

## license.capac events

### license.capac.enforce

#### Severity

ALERT

#### Description

This message occurs when the total storage capacity of all disks currently assigned to your system exceeds the limits of your capacity license. The system will shut down within the next 72 hours.

#### Corrective Action

To avoid disruption, either contact your sales representative immediately to obtain a license with a larger maximum capacity, or reduce the storage capacity of the system.

#### Syslog Message

"Storage capacity is at %llu GB, max allowed is %llu GB. Forced system shutdown is imminent."

#### Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**max\_capacity** (LONGINT): Maximum storage capacity supported under the current license, in gigabytes.

### license.capac.eval

#### Severity

ALERT

#### Description

This message occurs when a capacity license intended only for evaluation purposes is approaching its end date. This expiration might result in the forced shutdown of the system.

#### Corrective Action

To avoid any disruption, copy the contents of the volumes of this system to a new location. Contact your sales representative to purchase a version of the system.

#### Syslog Message

"The system have %llu GB installed. Evaluation license ends on %s."

#### Parameters

**current\_capacity** (LONGINT): Current storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

### license.capac.eval.ending

#### Severity

ERROR



## Description

This message occurs when a capacity license intended only for evaluation purposes is approaching its end date. It will soon not be possible to either create data aggregates, expand them, mirror them, or bring them online.

## Corrective Action

To avoid any disruption, copy the contents of the volumes of this system to a new location. Contact your sales representative to purchase a version of the system.

## Syslog Message

"The system is using %llu GB. Evaluation license ends on %s."

## Parameters

**current\_capacity** (LONGINT): Current storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

## license.capac.eval.enforce

### Severity

ALERT

### Description

This message occurs when the evaluation period has expired. Creating aggregates, extending them, mirroring them, and bringing them online will be prohibited. However, if the effective version of the cluster precedes the version of ONTAP® software on this node, any node with this condition might shut down within the next 72 hours.

### Corrective Action

To avoid any disruption, move or copy volumes as needed from this system. Contact your sales representative immediately to purchase a version of this system.

### Syslog Message

"Evaluation version of capacity license expired. Capacity at %llu GB. Certain aggregate operations are not allowed after %s."

### Parameters

**current\_capacity** (LONGINT): Current storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

## license.capac.eval.shutdown

### Severity

EMERGENCY

### Description

This message occurs when the evaluation period has expired. The system will be shut down within the next 10 minutes.

### Corrective Action

Contact your sales representative immediately to purchase a version of this system. Restart the system as needed to move or copy volumes.

## Syslog Message

"Storage capacity is at %llu GB. Evaluation license ends on %s. Forced system shutdown is imminent."

## Parameters

**current\_capacity** (LONGINT): Current storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

## license.capac.invalidated

### Severity

ERROR

### Description

This message occurs when the system determines that its existing capacity license does not cover its current used capacity during an attempted upgrade. To allow the upgrade to proceed, the system has deleted the license.

### Corrective Action

Either delete enough data aggregates to bring the system into compliance with the existing capacity license, or contact your sales representative immediately to obtain a new license with a sufficient capacity value. In either case, apply the license after the upgrade is successfully completed.

## Syslog Message

"Storage capacity now %llu GB, max allowed is %llu GB. Capacity license removed for upgrade."

## Parameters

**current\_capacity** (LONGINT): Current capacity of the system, in gigabytes.

**max\_capacity** (LONGINT): Maximum usable storage capacity supported under the current license, in gigabytes.

## license.capac.promo.failed

### Severity

ERROR

### Description

This message occurs when an evaluation version of ONTAP® Select has failed to be promoted to production.

### Corrective Action

Remove the license used during promotion by using the "system license clean-up -unused true" command, and then retry promotion.

## Syslog Message

"ONTAP Select promotion from evaluation to production was unsuccessful. Failure reason: %s."

## Parameters

**reason** (STRING): Reason for promotion failure.

## license.capac.promo.incomp

### Severity

ALERT

### Description

This message occurs when an evaluation version of ONTAP® Select has fails to automatically reboot to complete promotion.

### Corrective Action

Manually reboot the node to complete the evaluation to production promotion process.

### Syslog Message

"ONTAP Select promotion from evaluation to production failed to automatically reboot node to complete promotion. Node will use license (serial number: %s) after the node is manually rebooted."

### Parameters

**license\_serial\_number** (STRING): License serial number of the license intended to be used by the node after promotion.

## license.capac.promo.partial

### Severity

ALERT

### Description

This message occurs when an evaluation version of ONTAP® Select has partially completed the promotion process.

### Corrective Action

Manually reboot the node to complete the evaluation to production promotion process.

### Syslog Message

"ONTAP Select promotion from evaluation to production is incomplete. Node will use license (serial number: %s) after the node reboots."

### Parameters

**license\_serial\_number** (STRING): License serial number of the license intended to be used by the node after promotion.

## license.capac.promoted

### Severity

NOTICE

### Description

This message occurs when an evaluation version of ONTAP® Select has been successfully promoted to production.

### Corrective Action

(None).

## Syslog Message

"ONTAP Select promotion from evaluation to production was successful with a license (serial number: %s)"

## Parameters

**license\_serial\_number** (STRING): License serial number being used by the promoted node.

## license.capac.shutdown

### Severity

EMERGENCY

### Description

This message occurs when the total storage capacity of all disks currently assigned to your system has exceeded the limits of your capacity license. The system will shut down within the next 10 minutes.

### Corrective Action

Contact your sales representative immediately to obtain your capacity license. Restart the system as needed to complete the license installation.

## Syslog Message

"Storage capacity is at %llu GB, max allowed is %llu GB. Forced system shutdown is imminent."

## Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**max\_capacity** (LONGINT): Maximum storage capacity supported under the current license, in gigabytes.

## license.capac.unl.ending

### Severity

ERROR

### Description

This message occurs when a valid capacity license has not been installed. As a result, data aggregates cannot be created, expanded, mirrored, or brought online.

### Corrective Action

Contact your sales representative immediately to obtain your capacity license, to avoid any disruption.

## Syslog Message

"Now using %llu GB of storage capacity. Certain aggregate operations are prohibited after %s. The grace period for installing a valid capacity license is ending soon."

## Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**end\_date** (STRING): Date by which a valid license must be installed.

## license.capac.unl.enforce

### Severity

ALERT

## Description

This message occurs when a valid capacity license has not been installed. The system will not permit the creation, expansion, or mirroring of aggregates, or bringing aggregates online. However, if the effective version of the cluster precedes the version of ONTAP® on this node, any node with this condition may forcibly shut down.

## Corrective Action

Contact your sales representative immediately to obtain your capacity license, to avoid any disruption.

## Syslog Message

"The grace period for installing a valid capacity license has ended. Storage capacity is at %llu GB. Certain aggregate operations are prohibited as of %s."

## Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

## license.capac.unl.shutdown

### Severity

EMERGENCY

### Description

This message occurs when a valid capacity license has not been installed. The system will shut down within the next 10 minutes.

### Corrective Action

Contact your sales representative immediately to obtain your capacity license. Restart the system as needed to complete the license installation.

### Syslog Message

"Storage capacity is at %llu GB. License ends on %s. Forced system shutdown is imminent."

### Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**end\_date** (STRING): Evaluation period end date.

## license.capac.unlicensed

### Severity

ALERT

### Description

This message occurs when the system first detects that a valid capacity license has not been installed. As a result, data aggregates can be created, expanded, mirrored, and brought online only during the initial grace period. However, if the effective version of the cluster precedes the version of ONTAP® on this node, any node with this condition may be forcibly shut down after this grace period.

### Corrective Action

To avoid any disruption, contact your sales representative immediately to obtain your capacity license.

## Syslog Message

"No capacity license for system with %llu GB used. Certain aggregate operations are prohibited after %s."

## Parameters

**current\_capacity** (LONGINT): Current storage capacity of the system, in gigabytes.

**end\_date** (STRING): Unlicensed grace period end date.

## license.capac.valid

### Severity

NOTICE

### Description

This message occurs when a valid capacity license has been installed on your system.

### Corrective Action

(None).

## Syslog Message

"Capacity license allows %llu GB, you have %llu GB installed."

## Parameters

**max\_capacity** (LONGINT): Maximum storage capacity allowed under current license, in gigabytes.

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

## license.capac.warning

### Severity

NOTICE

### Description

This message occurs when the total storage capacity of all disks currently assigned to your system is in danger of exceeding the limits of your current license. Exceeding this capacity limit might result in the forced shutdown of the system.

### Corrective Action

To avoid disruption, either contact your sales representative immediately to obtain a license with a larger maximum capacity, or reduce the storage capacity of the system.

## Syslog Message

"Storage capacity is at %llu GB, max allowed is %llu GB."

## Parameters

**current\_capacity** (LONGINT): Currently used storage capacity of the system, in gigabytes.

**max\_capacity** (LONGINT): Maximum storage capacity supported under the current license, in gigabytes.

## license.cappool events

## license.cappool.aggr.noncomp

### Severity

ALERT

### Description

This message occurs when an aggregate does not comply with the current Capacity Pool entitlement, as determined by the Licensing Manager. In an HA configuration, noncompliant aggregates are prevented from being sent home during an HA takeover/giveback situation. In a non-HA configuration, these aggregates cannot be brought online if they are already offline.

### Corrective Action

Use the "system license license-manager check" command to determine whether the License Manager is accessible. Use the "system license show-aggregates -compliant false" command to find all aggregates that do not comply. From the web interface of the ONTAP® Select Deploy® instance, check the Capacity Pool licensing status, and ensure that sufficient capacity remains within the capacity pool used by this instance of ONTAP. If the capacity is insufficient, reduce the usage by instances of ONTAP that use the pool (for example, by deleting aggregates). If no reduction is possible, Contact NetApp technical support..

### Syslog Message

Aggregate %s (UUID: %s) is not compliant with current Capacity Pool entitlement.

### Parameters

**aggr\_name** (STRING): Name of the aggregate.

**aggr\_uuid** (STRING): UUID of the aggregate.

## license.cert events

### license.cert.expired

#### Severity

ALERT

#### Description

This message occurs when the license validation certificate is expired. No new license can be validated.

#### Corrective Action

Contact your sales representative or reseller about renewing your license validation certificate.

#### Syslog Message

No new license can be validated. Certificate expired on %s.

#### Parameters

**exp\_date** (STRING): Expiration date.

### license.cert.expired.check

#### Severity

NOTICE

## Description

This message occurs when the license validation certificate is near its expiration date. After the certificate expires, no new license can be validated.

## Corrective Action

Contact your sales representative or reseller about renewing your license validation certificate.

## Syslog Message

Certificate to validate licenses will expire on %s.

## Parameters

**exp\_date** (STRING): Expiration date.

# license.check events

## license.check.expired

### Severity

ERROR

### Description

This message occurs when the license checker determines that the Base license has expired. If this is the case, you cannot add nodes to the cluster

### Corrective Action

Install a permanent Base license on this cluster.

### Syslog Message

License Checking: Expired Base license (%s)

### Parameters

**error** (STRING): Interpreted error string from the license validation check.

## license.check.failed

### Severity

ERROR

### Description

This message occurs when the license checker has run and has been unable to validate the Base license due to an internal error. If there is not a valid Base license, then no additional nodes can be added to the cluster.

### Corrective Action

Install a valid Base license on this cluster.

### Syslog Message

License Checking: internal error validating Base license (%s)



**Parameters**

**error** (STRING): Interpreted error string from the license validation check.

**license.check.invalid****Severity**

ERROR

**Description**

This message occurs when the license checker determines that the Base license is incorrect in some way. If Base license is invalid, you cannot add nodes to the cluster.

**Corrective Action**

Install a valid Base license on this cluster.

**Syslog Message**

License Checking: The Base license is invalid (%s)

**Parameters**

**error** (STRING): Interpreted error string from the license validation check.

**license.check.mirror****Severity**

ERROR

**Description**

This message occurs when the license checker has run and has determined that at least one volume is being mirrored but there isn't a valid Mirror license installed.

**Corrective Action**

Install a valid Mirror license on this cluster.

**Syslog Message**

License Checking: Mirror License failure: (%s)

**Parameters**

**volume** (STRING): The first volume detected that requires a Mirror license

**license.check.missing****Severity**

ERROR

**Description**

This message occurs when the license checker determines that the Base license is missing. If this is the case, you cannot add nodes to the cluster.

**Corrective Action**

Reinstall the Base license.

## Syslog Message

License Checking: Missing Base license

## Parameters

(None).

## license.check.nodes

### Severity

ERROR

### Description

This message occurs when the license checker determines that the number of nodes in the cluster exceeds the number of nodes allowed by the Base license

### Corrective Action

Install a Base license with the correct node count, or remove nodes from the cluster using the unjoin command until the number of total nodes is equal to, or less than, the specified node count limit.

## Syslog Message

License Checking: Current node count (%d) exceeds the licensed node limit (%d)

## Parameters

**nodecount** (INT): Number of nodes detected in the cluster  
**nodelimit** (INT): Number of nodes allowed by the Base license

## license.check.nonodes

### Severity

ERROR

### Description

This message occurs when the license checker determines that it cannot retrieve the number of nodes in the cluster to compare against the Base license node count. If this is the case, then you cannot add nodes to the cluster.

### Corrective Action

Re-install the Base license so that the node count is correctly reflected. If that fails to fix the problem please contact Customer Support.

## Syslog Message

License Checking: Cannot to retrieve the number of nodes in the cluster.

## Parameters

(None).

## license.check.warn.expired

### Severity

ERROR

## Description

This message occurs when the license checker determines that a license has expired.

## Corrective Action

Install a permanent license on this cluster.

## Syslog Message

The license for package "%s" expired on %s.

## Parameters

**package** (STRING): Name of the license package.

**expiration** (STRING): Expiration date.

## license.check.warn.expiring

### Severity

ERROR

### Description

This message occurs when the license checker determines that a license is within a week of its expiration date.

### Corrective Action

Install a permanent license on this cluster.

### Syslog Message

The license for package "%s" expires on %s.

### Parameters

**package** (STRING): Name of the license package.

**expiration** (STRING): Expiration date.

## license.feats.events

### license.feats.delete.failed

#### Severity

ERROR

#### Description

This message occurs when the system cannot successfully delete a license for a feature. (This message is currently used only for FlexCache® licenses.)

#### Corrective Action

The feature might be vetoing deletion for specific, valid reasons (e.g., it is still in use), or it might be in a state that is inconsistent with deletion. Examine the error condition for more information. Contact NetApp technical support.

#### Syslog Message

License for feature %s and serial %s not deleted: %s.

## Parameters

**feature** (STRING): Name of the feature.  
**serial** (STRING): Serial number associated with the license.  
**reason** (STRING): Error condition encountered.

## license.feats.deleted

### Severity

NOTICE

### Description

This message occurs when the system deletes a license for a feature. (This message is currently used only for FlexCache® licenses.)

### Corrective Action

(None).

### Syslog Message

License deleted: feature %s, serial %s.

## Parameters

**feature** (STRING): Name of the feature.  
**serial** (STRING): Serial number associated with the license.

## license.feats.install.failed

### Severity

ERROR

### Description

This message occurs when the system cannot successfully install a license for a feature. (This message is currently used only for FlexCache® licenses.)

### Corrective Action

Check whether the feature is supported by the target instance of ONTAP® software. Check whether the target instance supports the change in the entitlements as specified by the license. Contact NetApp technical support.

### Syslog Message

License for feature %s and serial %s not installed: %s.

## Parameters

**feature** (STRING): Name of the feature.  
**serial** (STRING): Serial number associated with the license.  
**reason** (STRING): Error condition encountered.

## license.feats.installed

### Severity

NOTICE

## Description

This message occurs when the system successfully installs a license for a feature. (This message is currently used only for FlexCache® licenses.)

## Corrective Action

(None).

## Syslog Message

License installed for feature %s: serial: %s, %s.

## Parameters

**feature** (STRING): Name of the feature.

**serial** (STRING): Serial number associated with the license.

**attrs** (STRING): Attributes of interest for this license, listed as a colon-separated key-value pairs.

# license.fpool events

## license.fpool.changed

### Severity

NOTICE

### Description

This message occurs when the maximum storage capacity limit for FabricPool is changed.

### Corrective Action

(None).

### Syslog Message

"FabricPool entitlement: maximum capacity limit now allows %u TB."

### Parameters

**max\_capacity** (LONGINT): Current maximum limit supported under the current entitlement, in TB.

## license.fpool.denied

### Severity

NOTICE

### Description

This message occurs when an attempt to alter the maximum storage capacity limit for FabricPool is denied by FabricPool.

### Corrective Action

(None).

### Syslog Message

"FabricPool entitlement: maximum capacity limit not changed to %u from %u TB."

## Parameters

**requested\_max\_capacity** (LONGINT): Requested maximum capacity supported, in TB.  
**current\_max\_capacity** (LONGINT): Current maximum capacity limit supported under the current entitlement, in TB.

## license.fpool.no.cloud

### Severity

NOTICE

### Description

This message occurs when an attempt to alter the maximum storage capacity limit for FabricPool is denied, because there is no support for licensing FabricPool directly on instances of Cloud ONTAP(tm).

### Corrective Action

(None).

### Syslog Message

"FabricPool entitlement: maximum capacity limit not changed to %u TB - licensing FabricPool directly on instances of Cloud ONTAP not supported."

## Parameters

**requested\_max\_capacity** (LONGINT): Requested maximum capacity supported, in TB.

## license.fpool.removed

### Severity

NOTICE

### Description

This message occurs when the license for FabricPool usage is deleted.

### Corrective Action

(None).

### Syslog Message

"FabricPool entitlement deleted."

## Parameters

(None).

## license.fpool.unsupported

### Severity

NOTICE

### Description

This message occurs when an attempt to alter the maximum storage capacity limit for FabricPool is denied, because FabricPool is not supported on this instance of ONTAP(tm).

## Corrective Action

(None).

## Syslog Message

"FabricPool entitlement: maximum capacity limit not changed to %u TB - feature is not supported."

## Parameters

**requested\_max\_capacity** (LONGINT): Requested maximum capacity supported, in TB.

## license.fpool.wrong.cluster

### Severity

ERROR

### Description

This message occurs when an attempt to alter the maximum storage capacity limit for FabricPool is rejected. The license used in the attempt contains an incompatible cluster ID.

### Corrective Action

Determine the ID of the cluster by using the "cluster identity show" command. Within the license itself, determine the value of the "HostID" field. If the field exists, compare its value to the cluster ID. Confirm that these values match. If not, then the cluster is not the intended target of this license. Apply the license to the correct cluster if possible. For all other scenarios, contact NetApp technical support.

## Syslog Message

"FabricPool entitlement: maximum capacity limit not changed to %u TB; license applies to cluster with ID %s."

## Parameters

**requested\_max\_capacity** (LONGINT): Requested maximum capacity supported, in TB.

**cluster\_ID** (STRING): Cluster ID as specified in license.

## license.ontapi events

### license.ontapi.eval

### Severity

NOTICE

### Description

This message occurs when ONTAPI enablement evaluation license is installed successfully.

### Corrective Action

Contact your sales representative to get a license if you wish to continue using ONTAPI past this expiration date.

## Syslog Message

ONTAPI enablement evaluation license installed. Expiration: "%s".

## Parameters

**expiration\_date** (STRING): Expiration date of ONTAPI enablement evaluation license.

## license.ontapi.valid

### Severity

NOTICE

### Description

This message occurs when a valid ONTAPI enablement license has been installed on your system.

### Corrective Action

(None).

### Syslog Message

ONTAPI enablement license installed. Expiration: "%s".

## Parameters

**expiration\_date** (STRING): Expiration date of ONTAPI enablement license.

## license.pool events

### license.pool.sfo.gb.veto

### Severity

ERROR

### Description

This message occurs when a storage failover giveback is vetoed because the aggregate has no valid capacity lease from the capacity pool. This might be because the capacity pool license has expired, because there is not sufficient capacity in the pool, or because the License Manager is not accessible.

### Corrective Action

If the node could not check the compliance status of an aggregate, check that the License Manager is accessible by using the "system license license-manager check" command. If the License Manager is not accessible, verify that it is up and running in ONTAP Select Deploy, and that there are no network connectivity issues. If an aggregate has been found to be out of compliance, verify on ONTAP Select Deploy that the capacity pool license has not expired, and that there is enough capacity remaining for the aggregate. For all other reasons, Contact NetApp technical support.

### Syslog Message

SFO giveback vetoed. %s.

## Parameters

**reason** (STRING): Explanation for the veto.

## license.sfo events



## license.sfo.gb.veto

### Severity

ERROR

### Description

This message occurs when a storage failover giveback is vetoed because a select capacity license has not been installed and the license installation grace period has ended.

### Corrective Action

Install a valid select capacity license, and then reissue the "storage failover giveback" command. Use the "system license show-status" command to determine which node is missing the capacity license.

### Syslog Message

SFO giveback vetoed. Missing select capacity license.

### Parameters

(None).

## license.sfo.gb.veto.eval

### Severity

ERROR

### Description

This message occurs when a storage failover giveback is vetoed because the evaluation period for the ONTAP® instance has ended.

### Corrective Action

Contact your sales representative for further assistance.

### Syslog Message

SFO giveback vetoed. Evaluation period ended.

### Parameters

(None).

## license.subscription.events

### license.subscription.enforcement

### Severity

EMERGENCY

### Description

This message occurs when your Cloud ONTAP® license subscription has expired but may still be renewed.

### Corrective Action

Contact your sales representative to have your subscription reactivated.

## Syslog Message

%s

## Parameters

**expiration\_msg** (STRING): Expiration message.

## license.subscription.reminder

### Severity

NOTICE

### Description

This message occurs when your Cloud ONTAP® license subscription is going to expire soon.

### Corrective Action

Contact your sales representative to renew your subscription as soon as possible.

## Syslog Message

%s

## Parameters

**reminder\_msg** (STRING): Reminder message.

## license.subscription.warning

### Severity

ALERT

### Description

This message occurs when your Cloud ONTAP® license subscription is nearly expired.

### Corrective Action

Contact your sales representative to renew your subscription immediately, to avoid any disruption.

## Syslog Message

%s

## Parameters

**warning\_msg** (STRING): Warning message.

## license.term events

### license.term.enforcement

#### Severity

ALERT

#### Description

This message occurs when your feature license has expired. The feature might no longer function.

**Corrective Action**

Contact your sales representative to have your feature license reactivated.

**Syslog Message**

%s

**Parameters**

**enforcement\_msg** (STRING): Enforcement message.

**license.term.reminder****Severity**

NOTICE

**Description**

This message occurs when your feature license is going to expire soon.

**Corrective Action**

Contact your sales representative to renew your feature license as soon as possible.

**Syslog Message**

%s

**Parameters**

**reminder\_msg** (STRING): Reminder message.

**license.term.warning****Severity**

ERROR

**Description**

This message occurs when your feature license has expired. The feature might stop working soon.

**Corrective Action**

Contact your sales representative to renew your feature license immediately to avoid any disruption.

**Syslog Message**

%s

**Parameters**

**warning\_msg** (STRING): Warning message.

# lifp events

## lifp.net events

### lifp.net.bdmismatch

#### Severity

ERROR

#### Description

This message occurs while assigning network ports to broadcast domains, when LIF Placement subnet discovery finds a port that is assigned to a broadcast domain that is different from the broadcast domain to which other ports in the same physical broadcast domain are assigned.

#### Corrective Action

Ports assigned to broadcast-domains need to be corrected by the admin. The ports in one physical broadcast domain must be assigned to a single broadcast-domain.

#### Syslog Message

In IPspace %s, other ports in the same physical broadcast domain are assigned to broadcast-domain %s but %s is assigned to broadcast-domain %s.

#### Parameters

**ipspacename** (STRING): IPspace name.

**broadcastdomainOnOtherPorts** (STRING): Broadcast-domain name.

**port** (STRING): Port name.

**broadcastdomainOnCurrentPort** (STRING): Broadcast-domain name.

# Ildp events

## Ildp.receive events

### Ildp.receive.maximum.entries

#### Severity

NOTICE

#### Description

This message occurs when the LLDP (Link Layer Discovery Protocol) daemon running on the storage controller receives more unique neighbor advertisements than it can cache. The storage controller can cache 64 unique entries and accommodates new entries after that by removing the oldest entries first. This event is typically observed when connected to Ethernet networks that do not support LLDP broadcasts.

#### Corrective Action

If the storage system is connected to an Ethernet network that is forwarding LLDP advertisements, disable LLDP on the storage system or disable forwarding LLDP advertisements on the switch.

#### Syslog Message

LLDP maximum entries exceeded, oldest record replaced by new record received on port %s.

#### Parameters

**port\_name** (STRING): The network port that received the LLDP advertisement.

# Imgr events

## Imgr.aggr events

### Imgr.aggr.CA.locks.dropped

#### Severity

EMERGENCY

#### Description

This message occurs when locks established through CIFS Continuously Available (CA) shares on an aggregate are dropped after takeover, giveback, or aggregate relocation. This causes clients to see a service disruption on the aggregate.

#### Corrective Action

Confirm that the aggregate is online by using the 'aggr show -aggregate [name of the aggregate] -fields state' command. Then reopen files to reestablish CA locks. For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "Imgr.aggr.CA.locks.dropped" keyword.

#### Syslog Message

Continuously Available (CA) locks were dropped on aggregate %s because %s.

#### Parameters

**aggregate** (STRING): Name of the aggregate where CA locks were dropped.

**reason** (STRING): Reason for dropping CA locks, such as "remote lock synchronization is in progress".

## Imgr.badstructure events

### Imgr.badStructure.found

#### Severity

ERROR

#### Description

This event is issued whenever the multi-protocol lock manager finds a bad file pointer and turns off the logging of lock manipulation to preserve trace information. This information can be of use during core analysis.

#### Corrective Action

(None).

#### Syslog Message

bad file pointer found, preserving trace information

#### Parameters

(None).

# Imgr.ca events

## Imgr.CA.recons.mcc.veto

### Severity

ALERT

### Description

This message occurs when the system vetoes MetroCluster(tm) switchback or negotiated switchover because one of the volumes affected by the switch has CIFS Continuously Available (CA) locks that are still being reconstructed.

### Corrective Action

Retry the switchback or negotiated switchover after verifying that lock reconstruction is complete by using the 'debug locks reconstruction show' diagnostic privilege command.

### Syslog Message

Vetoed MetroCluster switchback or negotiated switchover because CA lock reconstruction is in progress.

### Parameters

(None).

# Imgr.cb events

## Imgr.cb.waiting.limit.hit

### Severity

ERROR

### Description

This message occurs when a lock manager waiting lock callback object cannot be allocated due to reaching the limit for waiting lock callback objects in the associated pool. It is possible there are too many clients holding locks on files where other clients want locks. Clients will have problems accessing or locking files through NAS protocols while at or above the waiting lock callback object limit.

### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

### Syslog Message

The number of waiting callback objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u locks %u/%u owners %u/%u

### Parameters

**cb\_waiting\_total** (INT): Number of allocated waiting callback objects in the pool.  
**cb\_waiting\_limit** (INT): Allocated callback waiting objects limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.

**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.

## Imgr.delegations events

### Imgr.delegations.limit.hit

#### Severity

ERROR

#### Description

This message occurs when a lock manager WAN FlexCache® delegation object cannot be allocated due to reaching the limit for WAN FlexCache delegation objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols using a FlexCache over a WAN to an origin volume on the node generating this message while at or above the WAN FlexCache delegation object limit.

#### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

#### Syslog Message

The number of WAN FlexCache delegation objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u locks %u/%u owners %u/%u

#### Parameters

**delegations\_total** (INT): Number of allocated WAN FlexCache delegation objects in the pool.  
**delegations\_limit** (INT): Allocated WAN FlexCache delegation objects limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.



# Imgr.files events

## Imgr.files.limit.hit

### Severity

ERROR

### Description

This message occurs when a lock manager file object cannot be allocated due to reaching the limit for file objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the file object limit.

### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

### Syslog Message

The number of file objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: hosts %u/%u locks %u/%u owners %u/%u.

### Parameters

**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.

# Imgr.gb events

## Imgr.gb.CA.locks.lost

### Severity

NOTICE

### Description

This message occurs when giveback veto is overridden. Locks established through CIFS continuously available (CA) shares on the aggregate being given back are dropped because the giveback occurs before lock sync completion.

### Corrective Action

(None).

## Syslog Message

Dropping partner aggregate locks during lock sync on aggregate %s.

## Parameters

**aggr** (STRING): Name of the aggregate for which CA locks are lost.

## Imgr.gb.CA.locks.may.lost

### Severity

ALERT

### Description

This message occurs when giveback veto is overridden. Locks established using CIFS continuously available (CA) shares on the aggregate being given back might be dropped, because giveback commit occurs simultaneously with lock sync completion.

### Corrective Action

(None).

## Syslog Message

Continuously Available (CA) locks may be lost on aggregate %s due to forced giveback commit occurring simultaneously with lock sync completion.

## Parameters

**aggr** (STRING): Name of the aggregate for which CA locks might be lost.

## Imgr.hostfile events

### Imgr.hostfile.cleared

### Severity

INFORMATIONAL

### Description

This message occurs when the volume lock manager hosts file is unusable. In such cases, the hosts files are cleared and reinitialized.

### Corrective Action

None.

## Syslog Message

Clearing lock manager file %s on volume %s due to error: %s.

## Parameters

**file\_name** (STRING): Name of the file that is being cleared.

**vol** (STRING): Name of the volume.

**error** (STRING): Error code.

## Imgr.hostfile.create.abort

### Severity

ERROR

### Description

This message occurs when the volume lock manager fails to create the hosts file.

### Corrective Action

None.

### Syslog Message

Creating lock manager file %s on volume %s failed due to error: %s.

### Parameters

**aggr** (STRING): The name of the volume.

**file\_name** (STRING): The name of the file that was being cleared.

**error** (STRING): Error code.

## Imgr.hostfile.init.failed

### Severity

ERROR

### Description

This message occurs when the volume lock manager hosts file initialization fails.

### Corrective Action

None.

### Syslog Message

Initializing lock manager file %s on volume %s failed due to error: %s.

### Parameters

**aggr** (STRING): The name of the volume.

**file\_name** (STRING): The name of the file that is being initialized.

**error** (STRING): Error code.

## Imgr.hosts events

### Imgr.hosts.limit.hit

#### Severity

ERROR

#### Description

This message occurs when a lock manager host object cannot be allocated due to reaching the limit for host objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the host object limit.

## Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

## Syslog Message

The number of host objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u locks %u/%u owners %u/%u.

## Parameters

**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.

# Imgr.lock events

## Imgr.lock.reclaim.aborted

### Severity

ERROR

### Description

This message occurs when the system aborts the volume lock reclaim process. A volume attempts to reclaim file locks when you bring the volume online after a reboot or after a takeover or giveback event. This error is usually caused by the volumes going offline during the initialization process.

### Corrective Action

None;When the volumes come back online the reclaim process will be restarted.

### Syslog Message

Aborting Lock manager grace period initialization on volume %s due to error: %s.

### Parameters

**aggr** (STRING): The name of the aggregate.  
**error** (STRING): Error code.

## Imgr.lock.type.counts

### Severity

NOTICE

## Description

This message provides additional information about the current count of lock manager objects when a lock manager pool limit is reached. The combination of the counts from a message indicating a limit has been reached and this message provide context for investigating and resolving the limit issues.

## Corrective Action

(None).

## Syslog Message

The number of lock manager object types in pool %s are: share locks %u byte locks %u waiting locks %u callback waiting %u pfs (clustered Data ONTAP FlexCache delegations) %u cached locks %u WAN FlexCache delegations %u

## Parameters

**pool\_name** (STRING): Name of the limit pool.

**share** (INT): Number of allocated share locks in the pool.

**byte** (INT): Number of allocated byte locks in the pool.

**waiting** (INT): Number of allocated waiting locks in the pool.

**cb\_waiting** (INT): Number of allocated cb\_waiting locks in the pool.

**pfs** (INT): Number of allocated clustered Data ONTAP® FlexCache® delegations in the pool.

**cached** (INT): Number of allocated clustered Data ONTAP FlexCache cached locks in the pool.

**delegations** (INT): Number of allocated WAN FlexCache delegations in the pool.

# Imgr.locks events

## Imgr.locks.limit.hit

### Severity

ERROR

### Description

This message occurs when a lock manager lock object cannot be allocated due to reaching the limit for lock objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the lock object limit.

### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

### Syslog Message

The number of lock objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u owners %u/%u.

### Parameters

**locks\_total** (INT): Number of allocated locks in the pool.

**locks\_limit** (INT): Allocated locks limit for the pool.

**pool\_name** (STRING): Name of the limit pool.

**files\_total** (INT): Number of allocated files in the pool.

**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.

## **Imgr.locks.sync.complete**

### **Severity**

NOTICE

### **Description**

This message occurs when synchronization of client file locks ("locks sync") with the partner node completes over the aggregates noted in the previous "Imgr.locks.sync.start" message. Lock sync occurs for locks that were established using CIFS continuously available (CA) shares.

### **Corrective Action**

(None).

### **Syslog Message**

Lock sync completed on %s.

### **Parameters**

**aggr** (STRING): Lock sync occurs in two cases: 1. Aggregate relocation 2. Giveback In the case of aggregate relocation, the aggregate name is provided. In the case of giveback, "aggr" indicates that lock sync is occurring over all aggregates on the node.

## **Imgr.locks.sync.start**

### **Severity**

NOTICE

### **Description**

This message occurs when synchronization of client file locks ("locks sync") with the partner node starts over one or more aggregates. Lock sync occurs for locks that were established using CIFS continuously available (CA) shares.

### **Corrective Action**

(None).

### **Syslog Message**

Lock sync started on %s.

### **Parameters**

**aggr** (STRING): Lock sync occurs in two cases: 1. Aggregate relocation 2. Giveback In the case of aggregate relocation, the aggregate name is provided. In the case of giveback, "aggr" indicates that lock sync is occurring over all aggregates on the node.

## **Imgr.locks.threshold.hit**

## Severity

ERROR

## Description

This message occurs when lock manager lock object has reached the threshold limit for lock objects in the associated pool. The system might become overloaded. Clients will have problems accessing files through NAS protocols while at or above the lock object limit.

## Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

## Syslog Message

The current number of lock objects %u, with maximum limit of %u has reached the threshold limit percentage of %u for pool "%s". The current number/limit of other lock manager objects are: files %u/%u, hosts %u/%u and owners %u/%u.

## Parameters

**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.  
**threshold\_percentage** (INT): Threshold locks limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.

# Imgr.lockstate events

## Imgr.lockstate.cleanup.fail

### Severity

ERROR

### Description

This message occurs when the system cannot clean up the lock state associated with the active file system. The stale lock state can lead to problems accessing files through NAS protocols.

### Corrective Action

Contact NetApp technical support for assistance.

### Syslog Message

Lock state cleanup failed for '%s' during '%s' on volume '%s%s%s' (DSID %llu) with error %d.

## Parameters

**lock\_type** (STRING): Type of locks being cleaned up.  
**cleanup\_reason** (STRING): Reason to clean up locks.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier for the volume.  
**volDsid** (LONGINT): Data set identifier (DSID) of the volume.  
**code** (INT): Error code returned while trying to clean up lock state.

# Imgr.mcc events

## Imgr.mcc.oplock.recall.fail

### Severity

ERROR

### Description

This message occurs when the system vetoes MetroCluster™ switchback or negotiated switchover because CIFS oplock recall fails. Some oplocks could not be recalled because of CPU overload, slow CIFS client response, or a large number of open files.

### Corrective Action

Recall oplocks for Vservers participating in negotiated switchover or switchback operations by using the command 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' for volumes or qtrees with oplocks enabled. Then retry the MetroCluster operation, and reenable oplocks after the operation is complete. If the error persists, then contact NetApp technical support for assistance.

### Syslog Message

Vetoed MetroCluster switchback or negotiated switchover because CIFS oplock recall failed with internal error (%d).

### Parameters

**error** (INT): Error number.

# Imgr.meta events

## Imgr.meta.corrupted

### Severity

ERROR

### Description

This event occurs when the system detects that the lock manager host metafile is corrupted.

### Corrective Action

It is possible that the corruption affected other parts of the file system. Run wafliron on the volume to verify file system consistency.



## Syslog Message

Corrupted lock state file in volume %s: reinitializing.

## Parameters

**path** (STRING): Volume name.

# Imgr.mismatch events

## Imgr.mismatch.file.objects

### Severity

NOTICE

### Description

This message occurs when the current file object of a lock does not match the file object found through the file handle. This might be caused by a SnapMirror® promote operation where the Master Data Set ID (MSID) is remapped.

### Corrective Action

None.

### Syslog Message

(None).

### Parameters

**volume** (STRING): Name of the volume, based on the file handle.

**lockid** (STRING): Lock Universal Unique Identifier (UUID) of the sharelock.

**caller** (STRING): WAFL® spinnp operation that encountered this event.

# Imgr.move events

## Imgr.move.dump.lock.counts

### Severity

NOTICE

### Description

This message occurs when a volume move operation fails the cutover phase, due to volume quiesce timer expiration, and not all locks were transferred before the failure.

### Corrective Action

Run the "vserver locks show -vserver <vserver> -volume <volume>" command to check the locks and the client IP addresses they are associated with. For NFS v4 locks, run instead the "vserver locks nfsv4 show -fields client-name" command in all nodes to get the same information. Reduce locks so that the number of locks remaining in this volume is fewer than the number scanned, and then attempt the volume move cutover again.

### Syslog Message

The volume move of volume with DSID %llu did not cut over due to volume quiesce timer expiration. The volume move has scanned %d out of %d locks in the volume.

## Parameters

**volDsid** (LONGINT): Data set identifier (DSID) of the volume.

**scanned\_locks** (INT): Number of locks in the volume that were scanned before the failure.

**total\_locks** (INT): Total number of locks in the volume.

# Imgr.notify events

## Imgr.notify.host.skip.write

### Severity

INFORMATIONAL

### Description

This message occurs when the volume lock manager cannot write to the notify hosts file due to an error, for example, insufficient space. In such cases, writing to the notify hosts file is skipped.

### Corrective Action

(None).

### Syslog Message

Skipped writing to notify hosts file %s on volume %s due to error: %s.

## Parameters

**file\_name** (STRING): Name of the file that writing was skipped on.

**vol** (STRING): Name of the volume.

**error** (STRING): Error code.

# Imgr.owners events

## Imgr.owners.limit.hit

### Severity

ERROR

### Description

This message occurs when a lock manager owner object cannot be allocated due to reaching the limit for owner objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the owner object limit.

### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

### Syslog Message

The number of owner objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u, locks %u/%u.

## Parameters

**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.

## Imgr.owners.threshold.hit

### Severity

ERROR

### Description

This message occurs when a lock manager owner object has reached the threshold limit for owner objects in the associated pool. The system might become overloaded. Clients will have problems accessing files through NAS protocols while at or above the owner object limit.

### Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object Imgr\_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr\_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr\_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

### Syslog Message

The current number of owner objects %u with maximum limit of %u has reached the threshold limit percentage of %u for pool "%s". The current number/limit of other lock manager objects are: files %u/%u, hosts %u/%u and locks %u/%u.

## Parameters

**owners\_total** (INT): Number of allocated owners in the pool.  
**owners\_limit** (INT): Allocated owners limit for the pool.  
**threshold\_percentage** (INT): Threshold locks limit for the pool.  
**pool\_name** (STRING): Name of the limit pool.  
**files\_total** (INT): Number of allocated files in the pool.  
**files\_limit** (INT): Allocated files limit for the pool.  
**hosts\_total** (INT): Number of allocated hosts in the pool.  
**hosts\_limit** (INT): Allocated hosts limit for the pool.  
**locks\_total** (INT): Number of allocated locks in the pool.  
**locks\_limit** (INT): Allocated locks limit for the pool.

## Imgr.partner events

### Imgr.partner.locks.outOfSync

## Severity

ERROR

## Description

This message occurs when a partner node's mirrored lock state goes out of sync with the actual lock state.

## Corrective Action

Under normal conditions, no action is necessary. There is an automatic synchronize process that will attempt to resynchronize the lock state. You can wait for this process to synchronize the lock state, or you can manually resynchronize the lock state by running the 'debug locks mirrored sync' command in diag mode, with the Vserver, volume, and file ID supplied in this message.

## Syslog Message

Mirrored lock state is out of sync for %s%s/%s (DSID %08x) for file %d Snapshot copy %d (mirroring status %d communication status %d)

## Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.

**dsid** (LONGINT): Data set identifier.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.

**mirroring\_status** (INT): Error status from the mirroring operation.

**csn\_status** (INT): Communication status from the mirroring operation.

## Imgr.partner.locks.resynched

## Severity

INFORMATIONAL

## Description

This message occurs when a partner node's mirrored lock state is resynchronized with the actual lock state after having been out of sync previously.

## Corrective Action

The mirrored lock state is now synchronized between the partner nodes. You can verify the mirror state of the locks by running the 'debug locks mirrored show' command in diag mode, with the Vserver, volume, and file ID supplied in this message.

## Syslog Message

Mirrored lock state is now in sync for %s/%s (DSID %08x) for file %d Snapshot copy %d.

## Parameters

**vol** (STRING): Volume name.

**volident** (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.

**dsid** (LONGINT): Data set identifier.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.

# Imgr.persistent events

## Imgr.persistent.open.limit.hit

### Severity

NOTICE

### Description

This message occurs when the number of persistent open lock count reaches the maximum supported limit.

### Corrective Action

Run the "vserver locks show-mirrored-open-count -detailed" command to check the persisted open lock count and try closing the unused files to bring this count down to the currently supported maximum persistent open locks.

### Syslog Message

The number of persisted open locks count %u has reached the limit %u.

### Parameters

**locks\_total** (INT): Number of persisted open locks.

**locks\_limit** (INT): Maximum supported persisted open locks.

# Imgr.precommit events

## Imgr.precommit.oplock.recall.fail

### Severity

ALERT

### Description

This message occurs when the system fails aggregate giveback or relocation because SMB oplock or NFSv4 delegation recall fails. Giveback or relocation cannot proceed until all oplocks are recalled successfully or the lock manager pre-commit phase is forced, in which case cached writes in clients might be lost.

### Corrective Action

There are two options for allowing aggregate giveback or relocation to proceed: (1) Force the sendhome pre-commit phase. For aggregate relocation, this can be done by specifying the '-override-destination-checks true' option. For giveback, this can be done by specifying the '-require-partner-waiting false' option. (2) Recall oplocks for Vservers on aggregates participating in the operation by using the 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' command for volumes or qtrees with oplocks enabled. Then retry the operation, and reenale oplocks after the operation is complete. If the error persists, contact NetApp technical support for assistance.

### Syslog Message

Aggregate "%s" failed giveback or relocation because lock manager oplock recall failed. Error: %s.

### Parameters

**aggrname** (STRING): Aggregate name.

**error** (STRING): Error details.

## Imgr.precommit.oplock.recall.forced

### Severity

NOTICE

### Description

This message occurs when the lock manager oplock recall fails during aggregate giveback or relocation, but the pre-commit phase has been forced to proceed. It is possible that cached writes in clients might have been lost as a result.

### Corrective Action

(None).

### Syslog Message

Lock manager oplock recall failed for aggregate "%s" during giveback or relocation. The failure abort was overridden. Error: %s.

### Parameters

**aggrname** (STRING): Aggregate name.

**error** (STRING): Error details.

## Imgr.precommit.oplock.recall.timeout

### Severity

ALERT

### Description

This message occurs when the lock manager oplock recall fails during aggregate giveback or relocation because the recalls could not be completed within the internal timeout window. The timeout typically can expire due to CPU overload, slow client response, or a large number of open files.

### Corrective Action

The timeout can be ignored by specifying the '-override-vetoes true' option in the giveback or relocation command.

### Syslog Message

Lock manager oplock recall failed to be completed for aggregate "%s" within %d seconds.

### Parameters

**aggrname** (STRING): Aggregate name.

**timeout** (INT): Timeout value.

## Imgr.pruning.events

### Imgr.pruning.delayed

### Severity

NOTICE

## Description

This message occurs when there is a delay in pruning and removing unused entries in the volume lock manager hosts file. The current hosts file contents remain; unused entries were not removed.

## Corrective Action

(None).

## Syslog Message

Pruning lock manager on volume %s for file %s delayed due to error: %s.

## Parameters

**volume** (STRING): Name of the volume.

**file\_name** (STRING): Name of the host file being pruned.

**error** (STRING): Error description.

# Imgr.recons events

## Imgr.recons.CA.lock.dropped

### Severity

ERROR

### Description

This message occurs when locks established through CIFS Continuously Available (CA) shares are dropped after a takeover, giveback, or aggregate relocation. This is because the multi-protocol lock manager finds bad mirrored lock state for CA locks and skips reconstructing them. This causes a client to see a service disruption on the file to which the CA locks belong.

### Corrective Action

Reopen the file to reestablish CA locks.

### Syslog Message

A Continuously Available (CA) lock for volume %s%s/%s (DSID %08x) file %d snapshot copy %d reconnect key %d is dropped because of %s.

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.

**dsid** (LONGINT): Data set identifier.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.

**reconnect\_key** (INT): Open identifier to uniquely identify a durable or persistent open on a volume

**reason** (STRING): Description of the reason for dropping a CA lock.

# Imgr.sendhome events

## Imgr.sendhome.veto.bootarg.override

### Severity

NOTICE

### Description

This message occurs during an SFO aggregate giveback or relocation. It implies that the automated nondisruptive update procedure will override any veto by the lock-manager, due to the presence of a bootarg.

### Corrective Action

(None).

### Syslog Message

Because bootarg.gb.override.imgr.veto is set to true, the ANDU procedure will override any lock-manager veto regarding the giveback of the "%s" aggregate.

### Parameters

**aggregate** (STRING): Name of the aggregate.

## Imgr.session events

### Imgr.session.failed

#### Severity

INFORMATIONAL

#### Description

This event occurs when a SpinNP session has failed, which causes session-scoped (e.g. CIFS) locks on the D-blade to be dropped. Only affects locks acquired using the failed session. This will cause clients to see a service disruption until they re-connect to the server.

#### Corrective Action

None — clients will automatically try to reconnect to the server, re-locking files as necessary.

#### Syslog Message

Dropped session-scoped locks from session %s:%s:%s due to failed session.

#### Parameters

**initiator** (STRING): The blade ID of the initiator.

**uniquifier** (STRING): The session uniquifier.

**target** (STRING): The blade ID of the target.

## Imgr.sf events

### Imgr.sf.up.ready

#### Severity

NOTICE



## Description

This message occurs when the lock manager non-disruption subsystem for CIFS continuously available (CA) shares notifies the high-availability (HA) module that transition into the up state is allowed. After the lock manager has sent this notification, the HA module will allow certain failover options to proceed such as aggregate relocation and giveback.

## Corrective Action

(None).

## Syslog Message

Lock manager allowed high availability module to transition to the up state for the following reason: %s.

## Parameters

**reason** (STRING): Reason for allowing transition to occur. Valid reasons include: 1. Lock-sync complete 2. Lock-sync timed out 3. CSM Invalid 4. Partner down 5. Mirroring disabled from bootarg 6. SMB3 disabled

# Imgr.sh events

## Imgr.sh.oplock.recall.fail

### Deprecated

This EMS was replaced by the new "Imgr.precommit.oplock.recall.\*" variants below and is no longer needed.

### Severity

ALERT

### Description

This message occurs when the system vetoes aggregate giveback or relocation because CIFS oplock recall fails. Some oplocks could not be recalled because of CPU overload, slow CIFS client response, or a large number of open files. Giveback or relocation cannot proceed until all oplocks are recalled successfully or the veto is overridden, in which case cached writes in clients might be lost.

### Corrective Action

Recall oplocks for Vservers on aggregates participating in the operation by using the command 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' for volumes or qtrees with oplocks enabled. Then retry the operation, and reenable oplocks after the operation is complete. If the error persists, contact NetApp technical support for assistance.

### Syslog Message

Aggregate "%s" vetoed giveback or relocation because CIFS oplock recall failed. Internal error: %s.

### Parameters

**aggrname** (STRING): Aggregate name.

**error** (STRING): Error details.

# Imgr.vol events

## **Imgr.vol.lock.reclaim.partial**

### **Severity**

NOTICE

### **Description**

This message occurs when a volume cannot reclaim certain file locks when the system attempts to bring it online. This can occur after a reboot, a takeover, a giveback or an MCC event.

### **Corrective Action**

(None).

### **Syslog Message**

Lock reclaims for volume %s are partial due to the following error: %s.

### **Parameters**

**vol** (STRING): The name of the volume.

**error** (STRING): Error code.

# logger events

## logger.usr events

### logger.usr

#### Severity

NOTICE

#### Description

This event is issued when the user runs the 'logger' command. The parameter indicates the message that is sought to be notified and is derived from the string passed in to the logger command.

#### Corrective Action

(None).

#### Syslog Message

%s

#### Parameters

**message** (STRING): Message passed by the user to the logger command.

# login events

## login.auth events

### login.auth.loginDenied

**Severity**

NOTICE

**Description**

This event is issued when a login attempt is refused due to authentication failure.

**Corrective Action**

Use a valid username/password combination to login.

**Syslog Message**

(None).

**Parameters**

**message** (STRING): Details about the failed login attempt.

# Ism events

## Ism.merge events

### Ism.merge.alog.problem

#### Severity

NOTICE

#### Description

This message occurs when a block store encounters an error while merging an append log. Normally these merges should recover, but it is important to keep track of exactly when and where the problem happened to help with performance tuning and recovery options.

#### Corrective Action

(None).

#### Syslog Message

Block store LSM "%u" (level %u) of keyspace "%u" on volume "%s" has encountered the following error while merging an append log: "%u : %s"

#### Parameters

**Ism\_id** (INT): LSM identifier.  
**Ism\_level** (INT): LSM level.  
**keyspace\_id** (INT): Keyspace identifier.  
**volume\_name** (STRING): Volume name.  
**error\_code** (INT): Error code.  
**error\_string** (STRING): Error string.

### Ism.merge.slog.problem

#### Severity

NOTICE

#### Description

This message occurs when a block store encounters an error while merging a sorted log. Normally these merges should recover, but it is important to keep track of exactly when and where the problem happened to help with performance tuning and recovery options.

#### Corrective Action

(None).

#### Syslog Message

Block store LSM "%u" (level %u) of keyspace "%u" on volume "%s" has encountered the following error while merging a sorted log: "%u : %s"

#### Parameters

**Ism\_id** (INT): LSM identifier.  
**Ism\_level** (INT): LSM level.  
**keyspace\_id** (INT): Keyspace identifier.

**volume\_name** (STRING): Volume name.  
**error\_code** (INT): Error code.  
**error\_string** (STRING): Error string.

## Ism.pal events

### Ism.pal.problem

#### Severity

NOTICE

#### Description

This message occurs when a block store PAL operation encounters an error. Normally these PAL operations should recover, but it is important to keep track of exactly when and where the problem happened to help with performance tuning and recovery options.

#### Corrective Action

(None).

#### Syslog Message

Block store PAL has encountered the following error while processing operation "%s": "%u"

#### Parameters

**pal\_op** (STRING): PAL operation being done.  
**error\_code** (INT): WAFL error code.

## Ism.rangedelete events

### Ism.rangeDelete.problem

#### Severity

NOTICE

#### Description

This message occurs when a block store encounters an error while performing a range delete operation. Normally these deletes should recover, but it is important to keep track of exactly when and where the problem happened to help with performance tuning and recovery options.

#### Corrective Action

(None).

#### Syslog Message

Block store LSM "%u" (level %u) of keyspace "%u" on volume "%s" has encountered the following error while performing a range delete operation: "%u : %s"

#### Parameters

**Ism\_id** (INT): LSM identifier.  
**Ism\_level** (INT): LSM level.  
**keyspace\_id** (INT): Keyspace identifier.  
**volume\_name** (STRING): Volume name.  
**error\_code** (INT): Error code.

**error\_string** (STRING): Error string.

# lun events

## lun.clone events

### LUN.clone.log.full

#### Severity

NOTICE

#### Description

This message occurs when lun clone creation fails because the limit on the maximum number of supported clone log entries (32705) is exceeded.

#### Corrective Action

Retry upon the completion of the pending clone operations for this volume.

#### Syslog Message

Logging the LUN clone creation has failed because the limit on the maximum number of supported clone log entries in the volume %s (DSID %llu) has been exceeded.

#### Parameters

**volume\_name** (STRING): Name of the volume for which the LUN clone creation has failed.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which the LUN clone creation has failed.

### LUN.clone.snapshot.deleted

#### Severity

NOTICE

#### Description

This message occurs when the snapshot used to clone a LUN is deleted before the completion of the cloning process. Data may be inaccessible for this LUN, depending on how much of the data was copied before the snapshot was deleted.

#### Corrective Action

The LUN should be destroyed to prevent I/O errors from occurring.

#### Syslog Message

A snapshot containing a file or LUN used to create the LUN clone %s (Inode number %d) in volume %s (DSID %llu) was deleted before the completion of the cloning process.

#### Parameters

**lun\_path** (STRING): Path for the LUN clone for which the snapshot was deleted.

**lun\_inode\_num** (INT): Inode number of the LUN clone for which the snapshot was deleted.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the containing volume.



# lun.copy events

## LUN.copy.dstPresent

### Severity

INFORMATIONAL

### Description

This message occurs when a LUN appears in the public namespace of the destination volume during a LUN copy operation.

### Corrective Action

(None).

### Syslog Message

LUN copy job with UUID %s created destination LUN (%s) successfully.

### Parameters

**job\_uuid** (STRING): UUID of the LUN copy job.

**dst\_path** (STRING): Destination path of the LUN copy job.

## LUN.copy.started

### Severity

INFORMATIONAL

### Description

This message occurs when a LUN copy operation is successfully started.

### Corrective Action

(None).

### Syslog Message

LUN copy started from (DSID:%llu)%s to (DSID:%llu)%s with job-uuid %s.

### Parameters

**src\_vol\_dsid** (LONGINT): Source volume Data Set ID (DSID).

**src\_lun** (STRING): Source LUN.

**dst\_vol\_dsid** (LONGINT): Destination volume Data Set ID (DSID).

**dst\_lun** (STRING): Destination LUN.

**UUID** (STRING): UUID identifying the copy operation.

# lun.create events

## LUN.create.nvfailed.volume

### Severity

ERROR

## Description

This message occurs when a LUN is created in a volume that is in the nvfailed state.

## Corrective Action

Clear the nvfailed state of the containing volume, take the appropriate nvfail recovery action to recover from the possible data loss, and bring the LUN online. Note that the actual recovery steps vary based on the application and the host.

## Syslog Message

LUN %s was created in a volume %s (DSID %llu) that is in the nvfailed state. It is offline and in the nvfailed state.

## Parameters

**lun\_path** (STRING): Path for the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

## LUN.create.resized.base

### Severity

NOTICE

### Description

This message occurs when creating a LUN from a base file causes a permanent size change to that base file.

### Corrective Action

(None).

### Syslog Message

The base file of the LUN at path %s in volume %s (DSID %llu) has been resized.

### Parameters

**lun\_path** (STRING): Path for the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

## lun.data events

### LUN.data.unavailable

#### Severity

ERROR

#### Description

This message occurs when data is unable to be read from a capacity tier.

#### Corrective Action

Troubleshoot capacity tier connectivity problems.

## Syslog Message

Unavailable error was received for the object at path %s (fileid:%u) in volume %s (DSID %llu).

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace.

**file\_id** (INT): Identifying number for the suspect inode.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

# lun.destroy events

## LUN.destroy

### Severity

NOTICE

### Description

This message occurs when a LUN is destroyed.

### Corrective Action

(None).

## Syslog Message

LUN %s, vol %s (DSID %llu) destroyed (UUID: %s).

## Parameters

**lun\_path** (STRING): Path for the LUN that was destroyed.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the LUN or a hyphen (-) if the UUID cannot be determined.

**object\_type** (STRING): Type of object under notification.

# lun.inaccessible events

## LUN.inaccessible.restore

### Severity

NOTICE

### Description

This message occurs when a LUN or NVMe namespace that is inaccessible due to restore is found in a volume. Such an object can contain inconsistent data and will not be accessible for I/O or management operations.

### Corrective Action

The object can be destroyed using the option force-fenced for the command 'lun destroy'. It can also be restored from a snapshot or backup.

## Syslog Message

The LUN or NVMe namespace %s in volume %s (DSID %llu) is inaccessible due to restore.

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

# lun.inconsistent events

## LUN.inconsistent.blocks

### Severity

ERROR

### Description

This message occurs when the block containing a LUN or NVMe namespace's attributes is found to be inconsistent.

### Corrective Action

Recover the LUN or NVMe namespace from a good backup. If the unrecoverable RAID is in a Snapshot(tm) copy, delete the Snapshot copy. Contact NetApp technical support for more options.

## Syslog Message

RAID error was received for the attributes of the object at path %s (fileid:%u) in volume %s (DSID %llu).

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace.

**file\_id** (INT): Identifying number for the suspect inode.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

## LUN.inconsistent.filesystem

### Severity

EMERGENCY

### Description

This message occurs when a LUN or NVMe namespace is found whose on-disk structure is inconsistent in the WAFL® file system. Such a LUN or NVMe namespace is not accessible for I/O or management operations.

### Corrective Action

Run waffliron on the containing aggregate, and then recover the LUN from a backup. Contact NetApp technical support for more options.

## Syslog Message

The on-disk structure of the object at path %s in volume %s (DSID %llu) is inconsistent in the WAFL file system.

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

# lun.inplace events

## LUN.inplace.dmpRstrLUNonNS

### Severity

ERROR

### Description

This message occurs when an in-place restore from a dump fails and creates a metadata inconsistency because the source object type does not match the existing destination object. A LUN has been restored over an existing NVMe namespace at the same path.

### Corrective Action

Unmap the NVMe namespace for which the mismatched restore has created a metadata inconsistency by using the "subsystem map remove" command. Remove the NVMe namespace for which the mismatched restore has created a metadata inconsistency by using the "force-fenced" option for the "lun delete" command. Retry the dump restore. If the original restore location was chosen in error, recover that deleted NVMe namespace.

### Syslog Message

An in-place restore from dump to an existing destination at path %s in volume %s (DSID %llu) failed due to a vdisk type mismatch.

## Parameters

**lun\_path** (STRING): Path of the existing destination for which the restore failed.  
**volume\_name** (STRING): Name of the destination volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the destination volume.

## LUN.inplace.dmpRstrNSonLUN

### Severity

ERROR

### Description

This message occurs when an in-place restore from a dump fails and creates a metadata inconsistency because the source object type does not match the existing destination object. An NVMe namespace has been restored over an existing LUN at the same path.

### Corrective Action

Unmap the LUN for which the mismatched restore has created a metadata inconsistency by using the "lun unmap" command. Remove the object for which the mismatched restore has created a metadata inconsistency by using the "force-fenced" option for the "lun delete" command. Retry the dump restore. If the original restore location was chosen in error, recover that deleted LUN.

## Syslog Message

An in-place restore from dump to an existing destination at path %s in volume %s (DSID %llu) failed due to a vdisk type mismatch.

## Parameters

**lun\_path** (STRING): Path of the existing destination for which restore failed.  
**volume\_name** (STRING): Name of the destination volume.  
**volume\_dsids** (LONGINT): Data set ID (DSID) of the destination volume.

## LUN.inplace.dumpRestoreFail

### Severity

ERROR

### Description

This message occurs when an in-place restore of a LUN from a dump fails because the source LUN is of a higher version than this Data ONTAP® release supports.

### Corrective Action

To recover the source LUN's data from the dump, first restore the LUN to a new path on the destination volume where it will be restored as a regular file with the entire source data intact. Then create a new and unmapped LUN from this file using the 'lun create -f' command. The original destination LUN is inaccessible for I/O or management operations and can be deleted using the 'force-fenced' option for the command 'lun destroy'.

## Syslog Message

An in-place restore from dump to an existing destination LUN at path %s in volume %s (DSID %llu) failed.

## Parameters

**lun\_path** (STRING): Path of the existing destination LUN for which restore failed.  
**volume\_name** (STRING): Name of the destination volume.  
**volume\_dsids** (LONGINT): Data set ID (DSID) of the destination volume.

## LUN.inplace.restore.failed

### Severity

ERROR

### Description

This message occurs when an in-place SFSR (Single File SnapRestore®) operation fails and might have left a LUN or NVMe namespace in an inconsistent state.

### Corrective Action

To recover the LUN, restore it to a new path. After the LUN is restored it needs to be mapped and hosts need to discover the new device. The LUN in the original path can be deleted.

## Syslog Message

An in-place SFSR (Single File SnapRestore®) failed at path %s in volume %s (DSID %llu) and might have left the LUN in an inconsistent state.

## Parameters

**lun\_path** (STRING): Path of the LUN or NVMe namespace for which the restore operation failed.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the containing volume.

# lun.metadb events

## LUN.metadb.version.mismatch

### Severity

ERROR

### Description

This message occurs when a version mismatch is detected for the metafile containing LUN metadata.

### Corrective Action

Rebuild the LUN metadata metafile for the volume by running the 'lun rescan -build-oovc' command.

### Syslog Message

A version mismatch was detected for the metafile containing LUN metadata metafile in volume %s (DSID %llu).

## Parameters

**volume\_name** (STRING): Name of the volume for which a version mismatch of the metafile containing LUN metadata was detected.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which a version mismatch of the metafile containing LUN metadata was detected.

# lun.metafile events

## lun.metafile.dirCreateFailed

### Severity

ERROR

### Description

lun metafile directory creation failed

### Corrective Action

Check available space on the volume and make sure that the volume is not full or close to full.

### Syslog Message

Couldn't create vdisk metafile directory %s.

## Parameters

**path** (STRING): path for the metafile directory that Data ONTAP failed to create

## lun.metafile.OOVC.corrupt

**Severity**

ERROR

**Description**

This message occurs when inconsistent blocks are read while reading or writing to the Out Of Volume Container (OOVC) metafile. The OOVC metafile contains certain LUN and NVMe namespace configuration and metadata for the objects in a given volume. SAN management operations might be impaired while the OOVC is in this state.

**Corrective Action**

Use the 'lun rescan -build-oovc' command to rebuild the OOVC metafiles for the affected volume.

**Syslog Message**

RAID error inconsistency was received for the OOVC metafile in volume %s (DSID %llu).

**Parameters**

**volume\_name** (STRING): Name of the volume for which the OOVC metafile was found to be inconsistent.

**volume\_dsids** (LONGINT): Data set ID (DSID) of the volume for which the OOVC metafile was found to be inconsistent.

**lun.metafile.VTOC.corrupt****Severity**

ERROR

**Description**

This message occurs when inconsistent blocks are read while reading or writing to the Vdisk Table of Contents (VTOC) metafile. The VTOC metafile contains the inventory of LUNs and NVMe namespaces in a given volume. SAN management operations might be impaired while the VTOC is in this state.

**Corrective Action**

Use the 'lun rescan -build-vtoc' command to rebuild the VTOC metafile for the affected volume.

**Syslog Message**

RAID error inconsistency was received for the VTOC metafile in volume %s (DSID %llu).

**Parameters**

**volume\_name** (STRING): Name of the volume for which the VTOC metafile was found to be inconsistent.

**volume\_dsids** (LONGINT): Data set ID (DSID) of the volume for which the VTOC metafile was found to be inconsistent.

**lun.move events****LUN.move.dstPresent****Severity**

INFORMATIONAL

**Description**

This message occurs when a LUN appears in the public namespace of the destination volume during a LUN move operation.



## Corrective Action

(None).

## Syslog Message

LUN move job with UUID %s created destination LUN (%s) successfully.

## Parameters

**job\_uuid** (STRING): UUID of the LUN move job.

**dst\_path** (STRING): Destination path of the LUN move job.

## LUN.move.started

### Severity

INFORMATIONAL

### Description

This message occurs when a LUN move operation is successfully started.

### Corrective Action

(None).

## Syslog Message

LUN move started from (DSID:%llu)%s to (DSID:%llu)%s with job-uuid %s.

## Parameters

**src\_vol\_dsid** (LONGINT): Source volume Data Set ID (DSID).

**src\_lun** (STRING): Source LUN.

**dst\_vol\_dsid** (LONGINT): Destination volume Data Set ID (DSID).

**dst\_lun** (STRING): Destination LUN.

**UUID** (STRING): UUID identifying the move operation.

## lun.no events

### LUN.no.memory.available

#### Severity

NOTICE

#### Description

This message occurs when an operation on LUNs or NVMe namespaces on the node owning the given volume fails due to the unavailability of Data ONTAP® memory.

#### Corrective Action

Retry the operation.

#### Syslog Message

An operation on LUNs or NVMe namespaces on the node owning the volume %s (DSID %llu) failed due to unavailability of Data ONTAP® memory.

## Parameters

**volume\_name** (STRING): Name of the volume for which the operation has failed.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which the operation has failed.

# lun.nvfail events

## LUN.nvfail.option.on

### Severity

NOTICE

### Description

This message occurs when the nvfail option for the named volume is turned on. The volume contains LUNs or NVMe namespaces and it is recommended that you keep this option on for such volumes.

### Corrective Action

(None).

### Syslog Message

The nvfail option for the volume %s (DSID %llu) has been turned on. The volume contains LUN or NVMe namespaces and it is recommended that this option be kept on for such volumes.

## Parameters

**volume\_name** (STRING): Name of the volume for which the nvfail option was turned on.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which the nvfail option has been turned on.

## LUN.nvfail.vol.proc.complete

### Severity

ERROR

### Description

This message occurs when an inconsistency is detected in the nvlog during boot or takeover, resulting in the initiation of nvfail processing and the LUNs in the volume being brought offline. This behavior is expected and is intended to prevent any FC/iSCSI data access to LUNs in a volume that might have lost the last few updates to the nvlog.

### Corrective Action

Verify that the LUN maps are valid. After the LUNs are mapped properly, use the 'lun online' command to bring the LUNs online to make them available to hosts. Host side recovery actions may be required to check the consistency of the data in the LUNs.

### Syslog Message

LUNs in volume %s (DSID %llu) have been brought offline because an inconsistency was detected in the nvlog during boot or takeover.

## Parameters

**volume\_name** (STRING): Name of the volume for which nvfail processing was completed.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which nvfail processing was completed.

## LUN.nvfail.vol.proc.failed

### Severity

EMERGENCY

### Description

This message occurs when nvfail processing for LUNs in a volume fails to complete. Some of the LUNs in the volume might not be offline. FC/iSCSI data access to the LUNs in the volume is prevented until the nvfail state of the volume is cleared.

### Corrective Action

For any LUNs that are offline, use the 'lun online' command to bring them online and make them available to hosts. Host side recovery actions might be required to check the consistency of the data in the LUNs.

### Syslog Message

An inconsistency was detected in the nvlog during boot or takeover, and some or all LUNs in volume %s (DSID %llu) might not have been brought offline.

### Parameters

**volume\_name** (STRING): Name of the volume for which LUN nvfail processing failed.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which LUN nvfail processing failed.

## LUN.nvfail.vol.proc.started

### Severity

ERROR

### Description

This message occurs when an inconsistency is detected in the nvlog during boot or takeover, resulting in the initiation of nvfail processing and the LUNs in the volume being brought offline. This behavior is expected and is intended to prevent any FC/iSCSI data access to LUNs in a volume that might have lost the last few updates to the nvlog.

### Corrective Action

Verify that the LUN maps are valid. After the LUNs are mapped properly, use the 'lun online' command to bring the LUNs online to make them available to hosts. Host side recovery actions may be required to check the consistency of the data in the LUNs.

### Syslog Message

LUNs in volume %s (DSID %llu) have been brought offline because an inconsistency was detected in the nvlog during boot or takeover.

### Parameters

**volume\_name** (STRING): Name of the volume for which nvfail processing was initiated.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which nvfail processing was initiated.

## lun.offline events

### LUN.offline

**Severity**

NOTICE

**Description**

This message occurs when a LUN is brought offline manually.

**Corrective Action**

(None).

**Syslog Message**

LUN %s, vol %s (DSID %llu) was brought offline (UUID: %s).

**Parameters**

**lun\_path** (STRING): Path for the LUN that was brought offline.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the LUN.

**object\_type** (STRING): Type of object under notification.

## lun.online events

**LUN.online****Severity**

NOTICE

**Description**

This message occurs when a LUN is brought online manually.

**Corrective Action**

(None).

**Syslog Message**

LUN %s, vol %s (DSID %llu) was brought online (UUID: %s).

**Parameters**

**lun\_path** (STRING): Path for the LUN that was brought online.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the LUN.

**object\_type** (STRING): Type of object under notification.

## lun.oovc events

**LUN.oovc.rebuild.done****Severity**

NOTICE

## Description

This message occurs when the rebuilding of the Out Of Volume Container (OOVC) file is complete.

## Corrective Action

(None).

## Syslog Message

The OOVC has been rebuilt for the volume %s (DSID %llu).

## Parameters

**volume\_name** (STRING): Name of the volume for which the OOVC file was rebuilt.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which the OOVC file was rebuilt.

## LUN.oovc.version.mismatch

### Severity

ERROR

### Description

This message occurs when a version mismatch is detected for the Out Of Volume Container (OOVC) metafile.

### Corrective Action

Rebuild the OOVC metafile for the volume by running the 'lun rescan -build-oovc' command.

### Syslog Message

A version mismatch has been detected for the OOVC metafile in volume %s (DSID %llu).

### Parameters

**volume\_name** (STRING): Name of the volume for which a version mismatch of the OOVC metafile has been detected.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which a version mismatch of the OOVC metafile has been detected.

## lun.op events

### LUN.op.nvfail.offline

#### Severity

ERROR

#### Description

This message occurs when a LUN is restored or a LUN clone is created in a volume that is currently in the nvfailed state. It can also occur if the source LUN for the operation is in the nvfailed state.

#### Corrective Action

If the containing volume is in the nvfailed state, clear the nvfailed state and bring the LUN online. Host side recovery actions might be required to check the consistency of the data in the LUN before the LUN is brought online. Note that the actual recovery steps vary based on the application and the host.

## Syslog Message

LUN %s in volume %s (DSID %llu) is offline and in the nvfailed state.

## Parameters

**lun\_path** (STRING): Path for the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsids** (LONGINT): Data set ID (DSID) of the containing volume.

# lun.out events

## LUN.out.of.space

### Severity

EMERGENCY

### Description

This message occurs when a LUN is brought offline by the SCSI layer because of a write failure due to lack of space.

### Corrective Action

Add space to the volume, and then bring the LUN online using the "lun online" command.

## Syslog Message

LUN %s, vol %s (DSID %llu) is out of space (UUID: %s).

## Parameters

**lun\_path** (STRING): Path for the LUN that was brought offline.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsids** (LONGINT): Data set ID (DSID) of the containing volume.  
**object\_uuid** (STRING): UUID of the LUN.  
**object\_type** (STRING): Type of object under notification.

# lun.outplace events

## LUN.outplace.dumpRestoreFail

### Severity

ERROR

### Description

This message occurs when an out-of-place restore of a LUN from dump fails because the source LUN is of a higher version than this Data ONTAP® release supports.

### Corrective Action

Restore creates a regular file at the destination path, with the same data as in the source LUN. To recover the source LUN's data from the dump, you can create a new and unmapped LUN from this file using the 'lun create -f' command.

## Syslog Message

An out-of-place restore of a LUN from dump for the destination path %s in volume %s (DSID %llu) failed.

## Parameters

**lun\_path** (STRING): Destination path for the failed LUN restore.  
**volume\_name** (STRING): Name of the destination volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the destination volume.

# lun.provisioning events

## LUN.provisioning.failed

### Severity

ERROR

### Description

This message occurs during the creation of a LUN or NVME namespace when the provisioning for its metadata fails due to insufficient space in the aggregate. Attempts to map or bind the object will fail until the aggregate has sufficient free space for the metadata.

### Corrective Action

Add more space to the aggregate. Use the command 'lun map' or 'lun bind' to allow provisioning of the LUN's metadata to be completed.

## Syslog Message

LUN or NVMe namespace %s in volume %s (DSID %llu) was created, but provisioning of its metadata failed due to insufficient space in the aggregate.

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace for which metadata provisioning failed.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

# lun.restore events

## LUN.restore.complete

### Severity

INFORMATIONAL

### Description

This message occurs when the SFSR (Single File SnapRestore®) for a LUN or NVMe namespace is complete.

### Corrective Action

(None).

## Syslog Message

LUN or NVMe namespace was restored at path %s in volume %s (DSID %llu).

## Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace that was restored.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

# lun.scsi events

## LUN.SCSI.reservation.removal.failed

### Severity

ALERT

### Description

This message occurs when the SCSI reservation for one or more LUNs in the secondary or primary Vserver cannot be deleted during a Vserver disaster recovery(DR) resynchronize operation. This might result in the host accessing stale SCSI reservation entries for the LUN during a subsequent Vserver DR activate operation.

### Corrective Action

Use the "lun persistent-reservation clear" command to remove the LUN's SCSI reservation.

### Syslog Message

SCSI reservation for LUN %s in volume %s (DSID %llu) could not be deleted.

## Parameters

**lun\_path** (STRING): Path for the LUN whose SCSI reservation could not be removed.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

# lun.serial events

## lun.serial.timeWentBack

### Severity

INFORMATIONAL

### Description

This event occurs when the system clock has been set back in time. It prints the current LUN serial number counter.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**counter** (LONGINT): Current value of LUN serial number counter.



# lun.space events

## LUN.space.resv.not.honored

### Severity

NOTICE

### Description

This message occurs when the space reservation settings for the named volume are not honored in its containing aggregate.

### Corrective Action

Make sure that the space guarantee settings for the volume are set to 'file' or 'volume'. If the volume guarantee is disabled, make sure that there is enough free space in the aggregate to enable the volume guarantee.

### Syslog Message

Space reservations in volume %s (DSID %llu) are not honored, either because the volume space guarantee is set to 'none' or it is disabled due to lack of space in its containing aggregate.

### Parameters

**volume\_name** (STRING): Name of the volume for which reservations are not honored.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which reservations are not honored.

# lun.transition events

## LUN.transition.fail.no.i2p

### Severity

ERROR

### Description

This message occurs when the transition of a 7-Mode LUN to clustered Data ONTAP® fails because the i2p database for the destination volume is not fully built.

### Corrective Action

Wait a few seconds and then issue the 'lun transition start' command to retry the transition of LUNs in the volume.

### Syslog Message

LUNs in volume %s (DSID %llu) could not be transitioned because the i2p database for the volume has not been fully built.

### Parameters

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

## LUN.transition.fail.no.space

**Severity**

ERROR

**Description**

This message occurs when the transition of a 7-Mode LUN to clustered Data ONTAP® fails due to lack of space in the destination volume or containing aggregate.

**Corrective Action**

Set the 'filesystem-size-fixed' option to false for the volume, add 1MB space for each LUN in the volume, and then issue the 'lun transition start' command to retry the transition.

**Syslog Message**

LUN %s in volume %s (DSID %llu) could not be transitioned due to lack of space in the volume.

**Parameters**

**lun\_path** (STRING): Path for the LUN that could not be transitioned.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**LUN.transition.failed****Severity**

ERROR

**Description**

This message occurs when the transition of a LUN to clustered Data ONTAP® fails.

**Corrective Action**

Determine the cause of failure and the specific corrective action by looking at other events issued for LUNs. Use the "lun transition start" command to retry the transition.

**Syslog Message**

LUN %s in volume %s (DSID %llu) could not be transitioned.

**Parameters**

**lun\_path** (STRING): Path for the LUN that could not be transitioned.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

**LUN.transition.incons.attrs****Severity**

ERROR

**Description**

This message occurs when a LUN is found whose 7-Mode on-disk structure is inconsistent in the WAFL® file system. This LUN cannot be transitioned and is not accessible for management.

**Corrective Action**

If the volume was being transitioned, destroy the LUN by using the "lun transition 7-mode delete" advanced-privilege command in clustered Data ONTAP®. Issue the command "lun transition 7-mode start"

to retry transition on the volume.

### **Syslog Message**

The 7-mode on-disk structure of the LUN at path %s in volume %s (DSID %llu) is inconsistent in the WAFL file system.

### **Parameters**

**lun\_path** (STRING): Path for the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

## **LUN.transition.no.type**

### **Severity**

NOTICE

### **Description**

This message occurs when Data ONTAP® encounters a LUN that is scheduled for transition and for which the multiprotocol type of the LUN cannot be determined. Data ONTAP substitutes a default value of "linux" for the type and calculates geometry values to match.

### **Corrective Action**

(None).

### **Syslog Message**

Multiprotocol type for LUN "%s" in volume "%s" (DSID: %llu) could not be determined. Once transitioned, this LUN will be assigned the default type "linux".

### **Parameters**

**lun\_path** (STRING): Path of the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data Set ID (DSID) of the containing volume.

## **LUN.transition.unsupp.clone**

### **Severity**

ERROR

### **Description**

This message occurs when transition of a LUN to clustered Data ONTAP® fails because the LUN is a Snapshot® copy-backed LUN clone. This LUN cannot be transitioned and is not accessible for management.

### **Corrective Action**

Destroy the LUN by using the "lun transition 7-mode delete" advanced-privilege command in clustered Data ONTAP. To make your original LUN a supported candidate for transition, split your Snapshot copy-backed LUN clone prior to transition.

### **Syslog Message**

LUN %s in volume %s (DSID %llu) could not be transitioned.

## Parameters

**lun\_path** (STRING): Path for the LUN that cannot be transitioned.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsids** (LONGINT): Data set ID (DSID) of the containing volume.

## LUN.transition.unsupp.geom

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® encounters a 7-Mode LUN that is scheduled for transition and for which the extent size or cylinder size is absent or unsupported. Data ONTAP substitutes reasonable geometry values upon transition.

### Corrective Action

(None).

### Syslog Message

LUN "%s" in volume "%s" (DSID: %llu) has a geometry that is not supported in clustered Data ONTAP. Once transitioned, this LUN will be assigned reasonable geometry values.

## Parameters

**lun\_path** (STRING): Path of the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsids** (LONGINT): Data Set ID (DSID) of the containing volume.

## LUN.transition.unsupported.type

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® encounters a 7-Mode LUN that is scheduled for transition and for which the multiprotocol type is unsupported in clustered Data ONTAP. Data ONTAP substitutes a default value of "linux" for the type upon transition and calculates geometry values to match.

### Corrective Action

(None).

### Syslog Message

LUN "%s" in volume "%s" (DSID: %llu) has a multiprotocol type that is not supported in clustered Data ONTAP. Once transitioned, this LUN will be assigned the default type of "linux".

## Parameters

**lun\_path** (STRING): Path of the LUN.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsids** (LONGINT): Data Set ID (DSID) of the containing volume.

# lun.vdisk events

## LUN.vdisk.cnt.exceeded

### Severity

ERROR

### Description

This message occurs when a volume contains more than the supported limit of LUNs in the volume. Performance degradation in volume management operations may occur as a result and it will not be possible to create more LUNs in this volume.

### Corrective Action

It is recommended that LUNs be moved to other volumes.

### Syslog Message

Volume %s (DSID %llu) contains more than the supported limit for LUNs (%d).

### Parameters

**volume\_name** (STRING): Name of the volume which contains more than the supported limit of LUNs.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume which contains more than the supported limit of LUNs.

**max\_create\_limit** (INT): The maximum number of LUNs that can be created in a volume.

## LUN.vdisk.prov.cnt.exceeded

### Severity

ERROR

### Description

This message occurs when provisioning of metadata for some LUNs fails because the limit for provisioning LUNs in the volume has been exceeded.

### Corrective Action

It is recommended that LUNs be moved to other volumes and then mapped using the 'lun map' command.

### Syslog Message

Provisioning of metadata for some LUNs failed because the limit on the maximum supported LUNs in the volume %s (DSID %llu) has been exceeded. These LUNs cannot be discovered by hosts and will not be accessible for I/O. Metadata cannot be provisioned for more than %d LUNs.

### Parameters

**volume\_name** (STRING): Name of the volume for which the LUN provisioning failed.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which the LUN provisioning failed.

**max\_prov\_limit** (INT): Maximum number of LUNs for which metadata can be provisioned in a volume.

# lun.vol events

## LUN.vol.no.space

### Severity

ERROR

### Description

This message occurs when an operation on a LUN or NVMe namespace fails due to insufficient space in the volume.

### Corrective Action

Add space to the volume and retry the operation.

### Syslog Message

An operation on %s in volume %s (DSID %llu) failed due to lack of space in the volume.

### Parameters

**lun\_path** (STRING): Path for the LUN or NVMe namespace.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the containing volume.

## LUN.vol.proc.fail.no.space

### Severity

ERROR

### Description

This message occurs when the processing for LUNs in the volume has failed due to lack of space in the volume or containing aggregate.

### Corrective Action

Check accompanying event messages to determine whether the volume or containing aggregate lacks space. Add space to the volume or aggregate containing the volume as necessary before retrying the operation. If the volume was being transitioned, issue the command 'lun transition start' to retry; otherwise issue the command 'lun rescan -sync' to complete processing of LUNs in the volume.

### Syslog Message

Processing for LUNs in volume %s (DSID %llu) failed due to lack of space.

### Parameters

**volume\_name** (STRING): Name of the volume for which the processing failed.

**volume\_dsids** (LONGINT): Data Set ID (DSID) of the volume for which the processing failed.

## lun.vtoc events

### LUN.vtoc.rebuild.done

### Severity

NOTICE

**Description**

This message occurs when the rebuilding of the Vdisk Table of Contents (VTOC) file is complete.

**Corrective Action**

(None).

**Syslog Message**

The VTOC has been rebuilt for the volume %s (DSID %llu).

**Parameters**

**volume\_name** (STRING): Name of the volume for which the VTOC file was rebuilt.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which the VTOC file was rebuilt.

**LUN.vtoc.version.mismatch****Severity**

ERROR

**Description**

This message occurs when a version mismatch is detected for the Vdisk Table of Contents (VTOC) metafile.

**Corrective Action**

Rebuild the VTOC metafile for the volume by running the 'lun rescan -build-vtoc' command.

**Syslog Message**

A version mismatch has been detected for the VTOC metafile in volume %s (DSID %llu).

**Parameters**

**volume\_name** (STRING): Name of the volume for which a version mismatch of the VTOC metafile has been detected.

**volume\_dsid** (LONGINT): Data Set ID (DSID) of the volume for which a version mismatch of the VTOC metafile has been detected.

# malloc events

## malloc.failure events

### malloc.failure

**Severity**

ERROR

**Description**

This event is generated if malloc is called unsafely.

**Corrective Action**

(None).

**Syslog Message**

%s %s %s %s %s %s %s %s %s %s %s %s

**Parameters**

- f2** (STRING): Third frame of backtrace.
- f3** (STRING): Fourth frame of backtrace.
- f4** (STRING): Fifth frame of backtrace.
- f5** (STRING): Sixth frame of backtrace.
- f6** (STRING): Seventh frame of backtrace.
- f7** (STRING): Eighth frame of backtrace.
- f8** (STRING): Ninth frame of backtrace.
- f9** (STRING): Tenth frame of backtrace.
- f10** (STRING): Eleventh frame of backtrace.
- f11** (STRING): Twelfth frame of backtrace.
- f12** (STRING): Thirteenth frame of backtrace.

### malloc.failure.detail

**Severity**

ERROR

**Description**

This event is generated if malloc is called unsafely.

**Corrective Action**

(None).

**Syslog Message**

%s %s

**Parameters**

- description** (STRING): Description of failure.
- buffers** (STRING): Details of buffers that were predirtied.



# master events

## master.cast events

### master.cast.cm.vote.against

#### Severity

NOTICE

#### Description

This message occurs when the master service of the associated node cannot connect to current cluster master. It tries to cast a vote against current cluster master in the global database. The current cluster master is demoted if more than half of the cluster nodes vote against it, and a new cluster master is promoted.

#### Corrective Action

(None).

#### Syslog Message

Service ID %u on node ID %u (UUID: %s) has voted against current cluster master with sequence number %u and service ID %u based on "%s". The current MVIP address connection state is "%s" with duration %s. The best connection state is "%s".

#### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**sequenceNumber** (INT): Sequence number of the cluster master.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

**reason** (STRING): Reason for casting the vote against current cluster master.

**currentState** (STRING): Current MVIP address connection state.

**currentStateDuration** (STRING): Duration of current MVIP address connection state.

**highestState** (STRING): Best MVIP address connection state reached since the node connected to current cluster master.

## master.couldnt events

### master.couldnt.cast.cm.vote

#### Severity

NOTICE

#### Description

This message occurs when the master service on the associated node is unable to cast a vote to update the connection status with current cluster master in global database since it has lost connection to the database server.

## Corrective Action

(None).

## Syslog Message

Service ID %u on node ID %u (UUID: %s) failed to cast a vote "%s" to cluster master with sequence number %u and service ID %u due to exception code: %s. The connection state is "%s" with duration %s. The best connection state is "%s".

## Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**voteToCast** (STRING): Vote which the master service failed to cast. It can be "true" or "false".

**sequenceNumber** (INT): Sequence number of current cluster master.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

**exceptionCode** (STRING): Exception code.

**currentState** (STRING): Current MVIP address connection state.

**currentStateDuration** (STRING): Duration of current MVIP address connection state.

**highestState** (STRING): Best MVIP address connection state reached since the node connected to current cluster master.

## master.couldnt.demote.cm

### Severity

NOTICE

### Description

This message occurs when the associated node tries to update the global database to demote current cluster master but the operation fails due to an exception. This could happen if another cluster node already made the same database change, or the connection to global database has been lost.

### Corrective Action

(None).

### Syslog Message

Service ID %u on node ID %u (UUID: %s) failed to update database to demote cluster master with service ID %u with exception code: %s.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

**exceptionCode** (STRING): Exception code.

## master.demote events

## master.demote.cm.committed

### Severity

NOTICE

### Description

This message occurs when the associated node successfully updates the global database to demote current cluster master. As a result, a new cluster is promoted.

### Corrective Action

(None).

### Syslog Message

Service ID %u on node ID %u (UUID: %s) successfully updated database to demote cluster master with sequence number %u and service ID %u.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**sequenceNumber** (INT): Sequence number of the cluster master last known by this node.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

## master.new events

### master.new.connection.state

#### Severity

NOTICE

#### Description

This message occurs when the connection state from the associated node to current cluster master has changed.

#### Corrective Action

(None).

#### Syslog Message

Service ID %u on node ID %u (UUID: %s) has a new MVIP address connection state "%s" to cluster master with service ID %u. The best connection state is "%s".

#### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**currentState** (STRING): Current management virtual IP (MVIP) address connection state.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

**highestState** (STRING): Best MVIP address connection state reached since the node connected to current cluster master.

# master.releasing events

## master.releasing.cm

### Severity

NOTICE

### Description

This message occurs when the node is demoted from being the cluster master.

### Corrective Action

(None).

### Syslog Message

Service ID %u on node ID %u (UUID: %s) releases control as the cluster master.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

# master.taking events

## master.taking.over.cm

### Severity

NOTICE

### Description

This message occurs when the node is promoted to the cluster master and takes control as the master.

### Corrective Action

(None).

### Syslog Message

Service ID %u on node ID %u (UUID: %s) takes over as the cluster master.

### Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

# master.try events

## master.try.demote.cm

### Severity

NOTICE

## Description

This message occurs when the associated node can not connect to current cluster master and tries to demote it by making associated global database change. The database change might fail if multiple cluster nodes try to do the same thing, and one of them has updated global database successfully.

## Corrective Action

(None).

## Syslog Message

Service ID %u on node ID %u (UUID: %s) tries to demote the cluster master with sequence number %u and service ID %u from the global database based on the votes (forVotes=%u, againstVotes=%u, staleVotes=%u).

## Parameters

**serviceID** (INT): Service ID that identifies the associated master service.

**nodeID** (INT): Node ID number of the associated node.

**nodeUuid** (STRING): Node UUID of the associated node.

**sequenceNumber** (INT): Sequence number of the cluster master last known by this node.

**lastKnownClusterMasterServiceID** (INT): Service ID number of the cluster master last known by this node.

**forVotes** (INT): Number of votes to keep current cluster master in the poll.

**againstVotes** (INT): Number of votes to demote current cluster master in the poll.

**staleVotes** (INT): Number of stale votes to update current cluster master in the poll.

# maxcid events

## maxcid.limit events

### maxCID.limit.exceeded

#### Severity

ALERT

#### Description

This message occurs when system exceeds the maximum number of open network data connections on the node.

#### Corrective Action

Balance the network load across all nodes within the cluster in order to avoid disruptions in service. Contact NetApp technical support.

#### Syslog Message

Number of active data connections is %d and system limit is %d.

#### Parameters

**activeCID** (INT): Number of open data connections on the node.

**maxCID** (INT): Maximum allowed limit of data connections per node.

## maxcid.threshold events

### maxCID.threshold.exceeded

#### Severity

ERROR

#### Description

This message occurs when the system is approaching the maximum number of open data connections supported on the node.

#### Corrective Action

Balance the network load across all nodes within the cluster.

#### Syslog Message

Number of active data connections is %d and maximum supported threshold is %d.

#### Parameters

**activeCID** (INT): Number of open data connections on the node.

**maxCIDThreshold** (INT): Maximum threshold of data connections supported per node.

# mc events

## mc.channel events

### mc.channel.create.failed

**Severity**

NOTICE

**Description**

This message occurs on failure in new channel creation in Mirror Cache.

**Corrective Action**

(None).

**Syslog Message**

Failure in Mirror Cache on new channel — Name: %s

**Parameters**

**channel** (STRING): Channel name

## mc.post events

### mc.post.replay.stats

**Severity**

NOTICE

**Description**

This message occurs post replay and conveys information regarding bdatas inserted, rejected, duplicate, hits and misses.

**Corrective Action**

(None).

**Syslog Message**

Mirror Cache post replay stats — Name: %s Inserted: %llu Rejected: %llu Duplicate: %llu Hits: %llu Misses: %llu.

**Parameters**

**name** (STRING): Aggr name

**inserted** (LONGINT): Bdatas inserted

**rejected** (LONGINT): Bdatas rejected

**duplicate** (LONGINT): Duplicate bdatas

**hits** (LONGINT): Bdata hits

**misses** (LONGINT): Bdatas misses

# mcc events

## mcc.auso events

### mcc.auso.trigFailed

#### Severity

EMERGENCY

#### Description

This message occurs when a trigger for automatic unplanned switchover fails to issue the operation.

#### Corrective Action

Verify the health of each node and that automatic unplanned switchover is enabled on both by using the "metrocluster node show" command. Use the "metrocluster switchover -forced-on-disaster true" command to switch over the DR-site storage manually.

#### Syslog Message

A trigger for Automatic Unplanned Switchover failed to issue the operation. Reason: %s.

#### Parameters

**probableCauseDescription** (STRING): Further description of the probable cause.

**alertId** (STRING): Identifier for this condition.

**alertingResource** (STRING): Unique name of the resource responsible for this condition.

**alertingResourceName** (STRING): Name of the resource responsible for this condition.

**possibleEffect** (STRING): Possible effect of the failure.

**probableCause** (STRING): Probable cause of the failure.

**additionalInfo** (STRING): Additional information about this failure.

### mcc.auso.triggered

#### Severity

NOTICE

#### Description

This message occurs when the node triggers an automatic unplanned switchover on detection of a failure of the DR partner node.

#### Corrective Action

(None).

#### Syslog Message

The node has triggered an Automatic Unplanned Switchover of the DR partner. Reason: %s. If Automatic Unplanned Switchover was triggered because of heartbeat loss, any ongoing dump on the partner may be aborted.

#### Parameters

**cause** (STRING): Cause for triggering automatic unplanned switchover.



# mcc.check events

## mcc.check.failed

### Severity

ERROR

### Description

This message occurs when the MetroCluster (tm) check job has failed to complete all checks for the system.

### Corrective Action

Run "metrocluster operation history show" command to check the reason for this failure. Make sure that all nodes in each of the clusters are healthy, and both clusters are in normal mode using the "metrocluster node show" and "metrocluster show" commands. Then retry the operation using the command "metrocluster check run". If the problem persists, contact NetApp technical support.

### Syslog Message

The MetroCluster (tm) check job has failed to complete all checks for the system in cluster %s.

### Parameters

**Cluster** (STRING): Cluster where the check was running.

# mcc.config events

## mcc.config.assignConfigFail

### Severity

ERROR

### Description

This message occurs when the MetroCluster(tm) local node configuration assignment to persistent storage failed.

### Corrective Action

Contact NetApp technical support for assistance.

### Syslog Message

(None).

### Parameters

**error** (STRING): Description of assign config error.

**nodeID** (STRING): Node UUID.

**dbladeID** (STRING): Dblade UUID.

## mcc.config.auso.enableIgnore

### Severity

ALERT

## Description

This message occurs when an attempt to enable automatic unplanned switchover fails because the capability has been manually disabled by the user or disabled because of a failed switchover or switchback operation.

## Corrective Action

If you want to keep automatic switchover enabled, enable it using the advanced mode command "metrocluster modify -automatic-switchover-onfailure true".

## Syslog Message

Failed to enable Automatic Unplanned Switchover because the capability has been manually disabled by the user or disabled because of a failed Switchover or Switchback operation.

## Parameters

(None).

## mcc.config.auso.stDisabled

### Severity

ERROR

### Description

This message occurs when automatic unplanned switchover capability is disabled.

### Corrective Action

Run the "metrocluster modify -node-name <nodename> -automatic-switchover-onfailure true" command for each node in the cluster to enable automatic switchover.

### Syslog Message

The state of Automatic Unplanned Switchover capability has been disabled.

### Parameters

**alertId** (STRING): Identifier for this condition.

**alertingResource** (STRING): Unique name of the resource responsible for this condition.

**alertingResourceName** (STRING): Display name of the resource responsible for this condition.

**possibleEffect** (STRING): Possible effect of the event.

**probableCause** (STRING): Probable cause of the event.

**probableCauseDescription** (STRING): Further description of the probable cause.

**additionalInfo** (STRING): Additional information about this event.

## mcc.config.auso.stEnabled

### Severity

NOTICE

### Description

This message occurs when automatic unplanned switchover capability is enabled.

### Corrective Action

(None).

## Syslog Message

The state of Automatic Unplanned Switchover capability has been enabled.

## Parameters

**alertId** (STRING): Identifier for this condition.

**alertingResource** (STRING): Unique name of the resource responsible for this condition.

## mcc.config.bootConfigNotRead

### Severity

ALERT

### Description

This message occurs when the MetroCluster(tm) configuration cannot be read from persistent storage. This could be because the partner has active reservations on the disk, the root aggregate is not assimilated, or the configuration has not been performed on the system yet. The system will try to retrieve or set the configuration try to retrieve or set the configuration automatically later during boot.

### Corrective Action

If the problem persists, contact NetApp technical support for assistance in resolving this issue.

## Syslog Message

Configuration could not be read from persistent storage due to internal error %s. The system will try to retrieve or set the configuration automatically later, during boot.

## Parameters

**error** (STRING): Actual error that caused configuration read to fail.

## mcc.config.checkPartnerError

### Severity

ERROR

### Description

This message occurs when an error is found during a MetroCluster(tm) configuration check. The checks include partner operational status and valid disaster recovery (DR) group configurations.

### Corrective Action

Make sure that all of the nodes are online and healthy by running the "cluster show" command and verifying that the "Health" and "Eligibility" fields displayed as true. Check the switch and controller configuration and make sure that all of the controllers in the DR group and HA/DR Interconnect ports are properly connected for DR group communication.

## Syslog Message

(None).

## Parameters

**partner\_names** (STRING): List of disaster recovery partner names.

**error** (STRING): Description of the error that occurred.

## **mcc.config.disableDrMirror**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when NVlog mirroring to the disaster recovery(DR) partners gets disabled persistently in a MetroCluster(tm) configuration.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**reason** (STRING): Reason the disaster recovery (DR) NVlog mirroring is disabled persistently.

## **mcc.config.disableMirrorMem**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when NVlog mirroring to the disaster recovery(DR) partners gets disabled in a MetroCluster(tm) configuration.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**reason** (STRING): Reason the DR NVlog mirroring is disabled in memory.

## **mcc.config.drPartner.mbxFail**

### **Severity**

ALERT

### **Description**

This message occurs when a node in a 2-node MetroCluster(tm) configuration fails to access its DR partner's mailbox. This may impact automatic unplanned switchover.

### **Corrective Action**

(None).

### **Syslog Message**

Read of DR partner's mailbox failed with error (%d).

**Parameters**

**err** (INT): Error returned from mailbox access.

**mcc.config.drPartner.mbxRead****Severity**

INFORMATIONAL

**Description**

This message occurs when a node in 2-node MetroCluster(tm) configuration is successfully able to access it's DR partner's mailbox.

**Corrective Action**

(None).

**Syslog Message**

Read of DR partner's mailbox successful.

**Parameters**

(None).

**mcc.config.enableDrMirror****Severity**

INFORMATIONAL

**Description**

This message occurs when NVlog mirroring to the disaster recovery(DR) partners gets persistently enabled in a MetroCluster(tm) configuration.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mcc.config.enableMirrorMem****Severity**

INFORMATIONAL

**Description**

This message occurs when NVlog mirroring to the disaster recovery(DR) partners gets enabled in a MetroCluster(tm) configuration.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mcc.config.headswapComplete****Severity**

NOTICE

**Description**

This message occurs when headswap update completes successfully on a node.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mcc.config.headswapMsgSent****Severity**

NOTICE

**Description**

This message occurs when a headswap HA message is successfully sent to a partner node.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**role** (INT): Role of the node.

**mcc.config.headswapNoNotify****Severity**

ERROR

**Description**

This message occurs when the MetroCluster(tm) node cannot communicate with one or more of its partners after a headswap. In a MetroCluster configuration, partners need to be informed of the old and new system ID of the local node through an HA message, so that they use the correct partner system ID during disaster recovery(DR) operations. Data ONTAP® takes appropriate recovery actions.

### Corrective Action

Other than making sure that all of the partners are online, there is no explicit user intervention required; the node automatically retries the operation at regularly scheduled intervals (every 30 seconds).

### Syslog Message

Local node system ID changed from %d to %d. Local node could not inform its HA or DR partner node of this change due to error %s.

### Parameters

**old\_system\_id** (INT): Old system ID of the node.

**new\_system\_id** (INT): New system ID of the node.

**error** (STRING): Internal error description identifying the reason for the failure.

## mcc.config.headswapPartnerEr

### Severity

ERROR

### Description

This message occurs when the partner returns a failure during a MetroCluster(tm) headswap configuration update. This can be due to configuration mismatches on the remote partner, disk ownership reassignment failure, or persistent storage write failure.

### Corrective Action

This could potentially be because the partner that failed was not fully functional or was operating in a degraded state. Make sure that the partner is online by issuing the "metrocluster show" and the "metrocluster node show" command. The partner configuration update will be retried at regular intervals automatically. In case of other configuration mismatch failure, check the relevant subsystem cited in the error and run the "disk show" command to examine the disk state. If the problem persists, contact NetApp technical support.

### Syslog Message

Partner %s failed to update configuration due to error %s.

### Parameters

**partner** (STRING): Name of the partner node.

**error** (STRING): Error description returned by the partner.

## mcc.config.headswapRevert

### Severity

NOTICE

### Description

This message occurs when the receiver of a headswap HA message detects that the old head was brought back in before headswap processing completed. The receiver will undo any nvram id update it had processed earlier.

### Corrective Action

(None).

**Syslog Message**

(None).

**Parameters**

**old\_sys\_id** (INT): Old system id received from the HA message

**new\_sys\_id** (INT): New system id received from the HA message

**mcc.config.locHeadswapDetect****Severity**

NOTICE

**Description**

This message occurs when a headswap is detected on the local node.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**old\_system\_id** (INT): Old system ID of the node.

**new\_system\_id** (INT): New system ID of the node.

**mcc.config.restoreFailure****Severity**

NOTICE

**Description**

Error returned while recovering the Configuration.

**Corrective Action**

(None).

**Syslog Message**

MetroCluster configuration was not restored from persistent source because of reason :%d.

**Parameters**

**err** (INT): Error returned while recovering the mailbox.

**mcc.config.restoreSuccessful****Severity**

INFORMATIONAL

**Description**

This message occurs when the MetroCluster(tm) configuration is successfully restored from an alternate persistent configuration source.



**Corrective Action**

(None).

**Syslog Message**

MetroCluster configuration was successfully restored from persistent source %s.

**Parameters**

**source** (STRING): Source from which the MetroCluster configuration was retrieved.

**mcc.config.selfNodeInvalid****Severity**

ERROR

**Description**

This message occurs when the local node UUID and the node UUID returned by the Connectivity, Liveness, Availability Monitor (CLAM) do not match. This is a configuration error and can result in failures during disaster recovery operations.

**Corrective Action**

Contact NetApp technical support for assistance.

**Syslog Message**

Configuration error: Local node UUID and the node UUID reported by CLAM do not match.

**Parameters**

**local\_node\_uuid** (STRING): Local node UUID.

**clam\_node\_uuid** (STRING): Node UUID returned by the CLAM.

**mcc.config.sysidNoNotify****Severity**

ERROR

**Description**

This message occurs when the local node is unable to notify its HA or DR partners of its system ID after a reboot.

**Corrective Action**

Verify that the HA or DR partner node is healthy and can be reached over the intercluster network.

**Syslog Message**

The node could not notify its nvram system id to partner %s. Error: %s.

**Parameters**

**partner** (STRING): Name of the partner node.

**error** (STRING): Error description.

**mcc.config.teardown**

**Severity**

INFORMATIONAL

**Description**

This message is generated whenever CTRL connections to MetroCluster(tm) partners are torn down.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**reason** (STRING): Reason for tearing down the CTRL connections.

**mcc.config.update.role****Severity**

INFORMATIONAL

**Description**

This message is generated whenever the roles of CTRL connections to MetroCluster(tm) partners are updated.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**reason** (STRING): Reason for updating the role of the CTRL connections.

**mcc.config.vso.disabled****Severity**

NOTICE

**Description**

This message occurs when the "vserver configuration override" command is disabled for MetroCluster over IP after controller replacement operations are completed.

**Corrective Action**

(None).

**Syslog Message**

MetroCluster over IP Vserver configuration override is disabled.

**Parameters**

(None).

## **mcc.config.vso.enabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the "vserver configuration override" command is enabled for MetroCluster over IP controller replacement operations.

### **Corrective Action**

(None).

### **Syslog Message**

MetroCluster over IP Vserver configuration override is enabled.

### **Parameters**

(None).

## **mcc.ctlbe events**

### **mcc.ctlbe.IODelayed**

#### **Severity**

INFORMATIONAL

#### **Description**

SCSI IO for RSM is delayed significantly in the iSCSI stack.

#### **Corrective Action**

(None).

#### **Syslog Message**

(None).

#### **Parameters**

**devSerial** (STRING): Device Serial Number

**lunId** (INT): The LUN ID.

**cdb** (STRING): A string containing the opcode byte of the CDB.

**delay** (LONGINT): The delay in milliseconds experienced by this IO in the software iSCSI stack.

### **mcc.ctlbe.IOResponseDelayed**

#### **Severity**

INFORMATIONAL

#### **Description**

SCSI IO for RSM is delayed significantly in the storage stack.

#### **Corrective Action**

(None).

## Syslog Message

(None).

## Parameters

**devSerial** (STRING): Device Serial Number

**lunId** (INT): The LUN ID.

**cdb** (STRING): A string containing the opcode byte of the CDB.

**delay** (LONGINT): The delay in milliseconds experienced by this IO in the D-blade storage stack.

# mcc.drdom events

## mcc.drdom.diskOwnerChange

### Severity

INFORMATIONAL

### Description

This message occurs when disk ownership changes are made by the disaster recovery disk ownership manager (DRDOM).

### Corrective Action

(None).

### Syslog Message

(%d) disks belonging to DR partner (%ld) underwent ownership change due to (%s).

### Parameters

**disk\_count** (INT): Number of disks that underwent ownership change.

**dr\_partner\_system\_id** (LONGINT): System ID of the DR partner whose disks underwent ownership change.

**reason** (STRING): Reason for the disk ownership change.

## mcc.drdom.diskReserveFail

### Severity

ERROR

### Description

This message occurs when one or more disks fails a disk reservation request from the disaster recovery disk ownership manager (DRDOM).

### Corrective Action

Make sure that none of the available DR partner disks are not in a failed or inconsistent state. Contact NetApp technical support if failed or misbehaving disks cannot be identified.

### Syslog Message

%d/%d DR disks failed disk reservation request from DRDOM, %d local disks.

### Parameters

**error\_count** (INT): Number of disks that failed disk reservation.

**disk\_num** (INT): Number of disks that have DR ownership.

**local\_disks** (INT): Number of disks that have local ownership.

## **mcc.drdom.disksAssignFail**

### **Severity**

ERROR

### **Description**

This message occurs when the disaster recovery disk ownership manager (DRDOM) fails to assign the ownership of one or more disks in the disk list enrolled in this operation to be reassigned.

### **Corrective Action**

Contact NetApp technical support for assistance.

### **Syslog Message**

DRDOM failed to assign the ownership of one or more disks in the disk list involved in disk ownership reassign request message during %s.

### **Parameters**

**operation** (STRING): The MetroCluster(tm) DR operation that issued the "assign disks" request.

**disks\_in\_request** (INT): Total number of disks for which the disk ownership change is requested.

**disks\_failed\_request** (INT): Total number of disks for which the disk ownership change request failed.

**error\_code** (INT): Internal error code for the failure.

## **mcc.drdom.hsAssignFail**

### **Severity**

ERROR

### **Description**

This message occurs when the node fails to change disk ownership for disaster recovery (DR) partner disks after a headswap. This can occur if the node is in switchover or it has disks that are left behind from previous DR operations.

### **Corrective Action**

Make sure that the DR partner disks with current-owner as local node are online and not in failed state or inconsistent state. Contact NetApp technical support if the failed or misbehaving disks cannot be identified.

### **Syslog Message**

Node failed to change disk ownership for DR partner disks after a headswap.

### **Parameters**

**disks\_in\_request** (INT): Total number of disks for which the disk ownership change is requested.

**disks\_failed\_request** (INT): Total number of disks for which the disk ownership change request failed.

**error\_code** (INT): Internal error code for the failure.

## **mcc.drgroup events**

### **mcc.drgroup.fetch.failed**

**Severity**

INFORMATIONAL

**Description**

This message occurs when MCC DR group information cannot be retrieved from RDB tables.

**Corrective Action**

(None).

**Syslog Message**

Unable to fetch DR-group information from RDB, error: %s.

**Parameters**

**reason** (STRING): Reason why RDB operation failed to fetch DR-group information.

## mcc.drmsg events

### mcc.drmsg.received

**Severity**

INFORMATIONAL

**Description**

This message occurs when a DR message is successfully received from a DR partner node.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**version** (INT): Version of the DR message.  
**type** (STRING): Type of message sent.  
**request** (STRING): Type of request being sent.  
**seq\_num** (LONGINT): Sequence number of the message.  
**rcv\_sysid** (LONGINT): System ID of receiving Node  
**sysid** (LONGINT): System ID of sending node.  
**rcv\_time** (LONGINT): Time when message was sent.  
**send\_time** (LONGINT): Time when message was sent.  
**status** (STRING): Status of the DR message.  
**fsm\_state** (STRING): FSM state of the responding node.  
**sb\_reason** (STRING): SB reason if any specified.  
**dr\_op\_state** (INT): DR operation state of the responding node.  
**ha\_fw\_state** (STRING): HA firmware state of the responding node.

### mcc.drmsg.sent

**Severity**

INFORMATIONAL

## Description

This message occurs when a DR message is successfully sent to a partner node.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**version** (INT): Version of the DR message.

**type** (STRING): Type of message sent.

**request** (STRING): Type of request being sent.

**seq\_num** (LONGINT): Sequence number of the message.

**sysid** (LONGINT): System ID of sending Node.

**send\_time** (LONGINT): Time when message was sent.

# mcc.drsom events

## mcc.drsom.disk.filter.fail

### Severity

INFORMATIONAL

### Description

This message occurs when it cannot be determined whether a disk belongs to a root aggregate.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk** (STRING): Name of the disk.

**op** (STRING): Operation during which this check failed.

**status** (INT): Status of the check that failed.

## mcc.drsom.disk.skipped

### Severity

INFORMATIONAL

### Description

This message occurs when disaster recovery switchover manager determines that a disk needs to be skipped from an assign request.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**disk** (STRING): Name of the disk.

**op** (STRING): Operation during which this check failed.

**reason** (STRING): Reason why the disk was skipped.

## mcc.drсом.fsm.execute

### Severity

INFORMATIONAL

### Description

This message occurs when the disaster recovery Switchover Monitor(DRSOM) FSM (Finite State Machine) starts processing an event.

### Corrective Action

(None).

## Syslog Message

DRSOM FSM state is (%s) and it will process the event (%s). Whether this event will be processed for the DR auxiliary partner: (%s)

## Parameters

**fsm\_state** (STRING): Current state of the DRSOM FSM.

**fsm\_event** (STRING): Event to be processed by the DRSOM FSM.

**DR\_aux** (STRING): Indication of whether the event will be processed for the DR auxiliary partner.

## mcc.drсом.fsm.invalid.state

### Severity

ALERT

### Description

This message occurs when the disaster recovery (DR) Switchover Monitor (DRSOM) FSM (Finite State Machine) reaches an invalid state.

### Corrective Action

If this problem is seen persistently, contact NetApp technical support for assistance in resolving this issue.

## Syslog Message

MCC DRSOM FSM state saw an invalid transition from %s to %s.

## Parameters

**prev\_st** (STRING): Previous DRSOM FSM state.

**st** (STRING): Current DRSOM FSM state.

## mcc.drсом.fsm.invalidState



**Severity**

ERROR

**Description**

This message occurs when an unexpected disaster recovery Switchover Monitor(DRSOM) FSM (Finite State Machine) event is generated.

**Corrective Action**

This can happen in cases where a new DR operation was issued before the previous one could finish and it might require a forced-recovery operation. Contact NetApp technical support for assistance in resolving this issue.

**Syslog Message**

DRSOM FSM event (%s) generated in state (%s) will cause an invalid FSM state transition.

**Parameters**

**fsm\_event** (STRING): DRSOM FSM event.

**fsm\_state** (STRING): DRSOM FSM state in which the event is generated.

**mcc.drsom.fsm.sb.rec.pending****Severity**

ALERT

**Description**

This message occurs when the disaster recovery switchover monitor (DRSOM) transitions to a switchback recovery pending state.

**Corrective Action**

If this node is booting up for the first time after a disaster, no action is required. If not, please check the states of all the nodes and contact NetApp technical support if they are incorrect.

**Syslog Message**

This node has entered the Switchback Recovery Pending State.

**Parameters**

(None).

**mcc.drsom.fsmState.EarlySb****Severity**

INFORMATIONAL

**Description**

This message occurs when the disaster recovery switchover monitor (DRSOM) recovers from unexpected error conditions by transitioning to perform an early switchback. An early switchback is performed in the following three cases: a node panics while processing a switchover; a node panics or reboots while retrying a switchover operation; or a node either encounters a subsystem failure, panics, or reboots while aborting a switchover operation. The switchover was reverted and the DRSOM state on all nodes is 'normal'.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mcc.drsom.mbxRecovery****Severity**

INFORMATIONAL

**Description**

This message occurs when the disaster recover Switchover Monitor(DRSOM) FSM (Finite State Machine) state stored persistently in the mailbox is inconsistent with the DRC (DR Coordinator) and has been recovered.

**Corrective Action**

(None).

**Syslog Message**

DRSOM FSM state stored in the mailbox is inconsistent with the DRC. It has been recovered and intialized to %s.

**Parameters**

**next\_fsm\_state** (STRING): State to which the DRSOM FSM will be initialized.

**mcc.drsom.process.op.fail****Severity**

ALERT

**Description**

This message occurs when the disaster recovery (DR) Switchover Monitor (DRSOM) cannot perform the DR operation.

**Corrective Action**

If this problem is seen persistently, contact NetApp technical support for assistance in resolving this issue.

**Syslog Message**

MCC state machine received invalid operation %s while it was in state %s.

**Parameters**

**op** (STRING): DR operation issued.

**st** (STRING): Current FSM (Finite State Machine) state.

**mcc.drsom.sbRecSt.change**

**Severity**

INFORMATIONAL

**Description**

This message occurs when there is a change in the disaster recovery Switchover Monitor(DRSOM) switchback recovery state.

**Corrective Action**

(None).

**Syslog Message**

Switchback recovery state changed from (%d) to (%d) due to (%s).

**Parameters**

**old\_state** (INT): Previous switchback recovery state.

**new\_state** (INT): New switchback recovery state.

**reason** (STRING): Reason for change of state.

**mcc.drsom.sstbl.configErr****Severity**

ERROR

**Description**

This message occurs during switchback processing, when the disaster recovery(DR) partner cannot be contacted to confirm that it is in a valid MetroCluster(tm) configuration state.

**Corrective Action**

Make sure that the DR partner is online and accessible by this node.

**Syslog Message**

(None).

**Parameters**

**partner** (STRING): Name of the failed partner.

**mcc.drsom.sstbl.sanownCompEr****Severity**

ERROR

**Description**

This message occurs during switchback processing, when the disaster recover(DR) partner cannot be contacted to enable or suspend a sanown completion event.

**Corrective Action**

Make sure that the DR partner is online and accessible by this node.

**Syslog Message**

(None).

## Parameters

**operation** (STRING): Disable or enable completion event.

**partner** (STRING): Name of the failed partner.

## mcc.drсом.sstblRunFailure

### Severity

ERROR

### Description

This message occurs when a particular subsystem cannot proceed with the disaster recovery (DR) operation.

### Corrective Action

Refer to the subsystem-specific error that would have been logged before this message.

### Syslog Message

(None).

## Parameters

**dr\_operation** (STRING): DR operation.

**sub\_system** (STRING): Name of the subsystem that encountered the failure.

## mcc.drсом.watch.conflict

### Severity

INFORMATIONAL

### Description

This message shows drсом watchdog information on conflicting operations

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**enabled** (INT): Current state of watchdog.

**extended** (INT): Current extended state of watchdog.

**last\_update\_time** (LONGINT): Last update time.

**current\_time** (LONGINT): Current time.

**last\_update\_count** (LONGINT): Last update count.

**last\_thread** (LONGINT): ID of last thread to update.

**last\_thread\_info** (STRING): Text describing last update.

**curr\_thread\_info** (STRING): Text describing current update.

## mcc.fabricpool events

## **mcc.fabricpool.license.absent**

### **Severity**

ERROR

### **Description**

This message occurs when a FabricPool license exists on only one of the two sites of a MetroCluster(tm) configuration.

### **Corrective Action**

Install a FabricPool license on the cluster that is missing the license.

### **Syslog Message**

The FabricPool license is installed on only one site of this MetroCluster configuration. The license is missing on cluster %s (ID: %s).

### **Parameters**

**cluster\_name** (STRING): Name of the cluster which is missing the license.

**cluster\_uuid** (STRING): UUID of the cluster.

## **mcc.fabricpool.license.diff**

### **Severity**

ERROR

### **Description**

This message occurs when the FabricPool license capacity or term limits are different for the two sites of a MetroCluster(tm) configuration.

### **Corrective Action**

Install a new license on one of the clusters so that the license limits match.

### **Syslog Message**

The FabricPool license limits are different for the two sites of this MetroCluster configuration. The license limits are %s and %s on site %s and %s and %s on site %s.

### **Parameters**

**local\_license\_limit** (STRING): License capacity limit from the local cluster.

**local\_term\_limit** (STRING): License term limit from the local cluster.

**local\_cluster\_name** (STRING): Name of the local cluster.

**remote\_license\_limit** (STRING): License capacity limit from the remote cluster.

**remote\_term\_limit** (STRING): License term limit from the remote cluster.

**remote\_cluster\_name** (STRING): Name of the remote cluster.

## **mcc.fabricpool.mirror.missing**

### **Severity**

ALERT

### **Description**

This message occurs when a FabricPool in a MetroCluster configuration does not have a mirror object store attached.

## Corrective Action

Use the "storage aggregate object-store mirror" command to add a mirror object store to the FabricPool.

## Syslog Message

FabricPool: "%s" UUID: "%s" in a MetroCluster configuration does not have synchronous data protection enabled because the mirror object store is missing.

## Parameters

**aggr** (STRING): Name of the FabricPool.

**UUID** (STRING): UUID of the FabricPool.

# mcc.forced events

## mcc.forced.auso.issued

### Severity

NOTICE

### Description

This message occurs when the node triggers a forced automatic unplanned switchover because the "allow-auto-forced-switchover" option is true.

### Corrective Action

(None).

### Syslog Message

The node has triggered a forced automatic unplanned switchover of the DR partner. Reason: "%s".

### Parameters

**reason** (STRING): Reason for triggering a forced automatic unplanned switchover with "allow-auto-forced-switchover" option.

# mcc.hwassist events

## mcc.hwassist.trap.AUSO

### Severity

NOTICE

### Description

This message occurs when a HW assist trap is received from the Service Processor of the DR-partner and is used to trigger an Automatic SwitchOver.

### Corrective Action

(None).

### Syslog Message

HW assist trap of type=%s, severity=%s received for event=%s from sysid=%s.

## Parameters

**type** (STRING): Type of event, such as `system_down`.  
**severity** (STRING): Severity of the event, such as `alert`.  
**event** (STRING): Specific event, such as `power_loss`.  
**sysid** (STRING): System-id of the node generating the event.

# mcc.ic events

## mcc.ic.process.node.adv

### Severity

INFORMATIONAL

### Description

This message occurs when the MetroCluster(tm) config module receives a node advertisement from any of its partners.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**sys\_id** (LONGINT): System id of the node sending the advertisements.  
**is\_local\_node\_assigned** (INT): Boolean specifying whether the local node has joined the cluster.  
**is\_partner\_assigned** (INT): Boolean specifying whether the node sending advertisements has joined the cluster.

# mcc.ipspace events

## mcc.ipspace.not.found

### Severity

ERROR

### Description

This message occurs when the IPspace is not found in a MetroCluster(tm) configuration while replicating the object store configuration. Until this is resolved, some MetroCluster and FabricPool operations will fail.

### Corrective Action

Create the missing MetroCluster IPspace. The configuration replication task will be completed automatically.

### Syslog Message

Object store "%s" configuration replication failed because IPspace "%s" was not found in cluster "%s".

## Parameters

**object\_store\_name** (STRING): Name of the object store.  
**ipspace\_name** (STRING): Name of the IPspace.

**cluster\_name** (STRING): Name of the cluster.

## mcc.kill events

### mcc.kill.received

#### Severity

INFORMATIONAL

#### Description

This message occurs when the node is halting because it received a "kill" message from the disaster recovery (DR) partner.

#### Corrective Action

(None).

#### Syslog Message

Node is halting because a kill message with reason %s has been received from the DR partner with system ID %ld.

#### Parameters

**reason** (STRING): Reason for halting the node.

**sender\_sysid** (LONGINT): System ID of the sender.

## mcc.lunmgr events

### mcc.lunmgr.lunid\_alloc\_failure

#### Severity

ERROR

#### Description

This message occurs when failing to allocate a CTL Target LUN ID to a target disk

#### Corrective Action

This indicates an internal error during the discovery of Target LUN devices.

#### Syslog Message

Device S/N %s - Failed to allocate a LUN ID

#### Parameters

**device\_serial\_number** (STRING): Device serial number

## mcc.network events

### mcc.network.congestion

#### Severity

NOTICE



## Description

This message occurs when a MetroCluster over IP configuration detects network congestion over the inter-cluster links (ISL) that connect sites and an action is taken to reduce further network congestion.

## Corrective Action

Verify that network between MetroCluster sites is following minimum SLA requirements and fix any such violations if needed. If the issue persists, contact NetApp technical support.

## Syslog Message

Network congestion detected. Action taken: %s.

## Parameters

**action\_taken** (STRING): Action taken to reduce network congestion.

# mcc.notify events

## mcc.notify.auto\_heal\_queue\_failure

### Severity

ERROR

### Description

Failed to queue the "MetroCluster Auto Heal" job

### Corrective Action

Manually complete the healing phase using the command "metrocluster heal aggregates" and "metrocluster heal root-aggregates"

### Syslog Message

(None).

### Parameters

**retries\_attempted** (INT): Number of retries attempted

# mcc.nso events

## mcc.nso.abort.dump

### Severity

NOTICE

### Description

This message occurs when dump/restore is aborted during a MetroCluster switch-over operation.

### Corrective Action

(None).

### Syslog Message

Dump/restore operations were aborted because %s is in progress.

## Parameters

**nso** (STRING): Describes whether a MetroCluster switch-over operation is progress.

## mcc.nso.veto.repl

### Severity

NOTICE

### Description

This message occurs when an MetroCluster(tm) switch-over is vetoed by a critical replication transfer (for example, volume move in the cutover phase).

### Corrective Action

Either abort the replication transfer that is preventing the MetroCluster switch-over from being completed, retry the MetroCluster switch-over operation with the '-override-vetoes true' option, or wait for the transfer to finish.

### Syslog Message

Could not complete switch-over because a replication transfer with UUID %s is in progress.

## Parameters

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed MetroCluster switch-over.

## mcc.nvmeof events

### mcc.nvmeof.port.modify

#### Severity

NOTICE

#### Description

This message occurs when the mode of an NVMe-oF port is changed.

#### Corrective Action

(None).

#### Syslog Message

Changing NVMe-oF port %s to %s mode.

#### Parameters

**port** (STRING): Name of port being configured.

**mode** (STRING): Mode in which to configure port: network or storage.

### mcc.nvmeof.rdma.disconnect.event

#### Deprecated

Deprecated as of version 9.13.1.

#### Severity

NOTICE

## Description

This message occurs when the NVMe-oF driver detects an RDMA disconnect. There might be a decrease in performance while reconnecting.

## Corrective Action

(None).

## Syslog Message

NVMe-oF RDMA disconnect on subnqn %s, controller ID %d, qpair ID %d.

## Parameters

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

## mcc.nvmeof.rdma.disconnect.status

### Deprecated

Deprecated as of version 9.13.1.

### Severity

NOTICE

## Description

This message occurs when the NVMeoF driver detects a command that finishes with a disconnect status. There might be a decrease in performance while reconnecting.

## Corrective Action

(None).

## Syslog Message

NVMe-oF disconnect on subnqn %s, controller ID %d, qpair ID %d.

## Parameters

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

## mcc.nvmeof.subsystem.add

### Severity

NOTICE

## Description

This message occurs when an NVMe-oF subsystem is added.

## Corrective Action

(None).

## Syslog Message

NVMe-oF subsystem added at address %s.

## Parameters

**address** (STRING): IP address of subsystem.

## mcc.nvmeof.subsystem.delete

## Severity

NOTICE

## Description

This message occurs when an NVMe-oF subsystem is deleted.

## Corrective Action

(None).

## Syslog Message

NVMe-oF subsystem deleted at address %s.

## Parameters

**address** (STRING): IP address of subsystem.

## mcc.nvmeof.timeout

## Severity

NOTICE

## Description

This message occurs when the NVMeoF driver detects a timeout on a request. There might be a decrease in performance while retrying.

## Corrective Action

(None).

## Syslog Message

Timeout on subnqn %s, controller ID %d, qpair ID %d, sequence number %d.

## Parameters

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

**cid** (INT): NVMe-oF CID.

## mcc.sb events

### mcc.sb.abort.dump

## Severity

NOTICE

**Description**

This message occurs when dump/restore is aborted during a MetroCluster switchback operation.

**Corrective Action**

(None).

**Syslog Message**

Dump/restore operations were aborted because %s is in progress.

**Parameters**

**sb** (STRING): Describes whether a MetroCluster switchback operation is progress.

**mcc.sb.veto.repl****Severity**

NOTICE

**Description**

This message occurs when an MetroCluster(tm) switchback is vetoed by a critical replication transfer (for example, volume move in the cutover phase).

**Corrective Action**

Either abort the replication transfer that is preventing the MetroCluster switchback from being completed, retry the MetroCluster switchback operation with the '-override-vetoes true' option, or wait for the transfer to finish.

**Syslog Message**

Could not complete switchback because a replication transfer with UUID %s is in progress.

**Parameters**

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed the MetroCluster switchback.

**mcc.sbca events****mcc.sbca.mdvNotAvailable****Severity**

ERROR

**Description**

This message occurs when a MetroCluster (tm) system metadata volume that provides storage for the configuration replication services is not available.

**Corrective Action**

Contact technical support for assistance.

**Syslog Message**

A MetroCluster (tm) system metadata volume is not available in cluster %s.

## Parameters

**Cluster** (STRING): Cluster attempting to use the volume.

# mcc.scsipd events

## mcc.scsipd.iScsiSessionReconnectDestroySim

### Severity

NOTICE

### Description

iSCSI session maintenance thread was signaled to reconnect and destroy SIM paths.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**watchdogIdx** (INT): The watchdog thread index for the mcc\_scsipd driver.

**watchdogTimeLimit** (INT): The time limit in seconds that the watchdog thread waits before signaling the iSCSI session maintenance thread.

# mcc.so events

## mcc.so.nodeObjectLmt.chkFail

### Severity

ALERT

### Description

This message occurs when a MetroCluster(tm) switchover operation cannot be performed on this node because it is in the takeover state and one of the nodes in its disaster recovery(DR) group has the "node-object-limit" MetroCluster configuration option set to off.

### Corrective Action

Use the "storage failover giveback" command to take this node out of takeover state.

### Syslog Message

Switchover operation cannot be performed on this node because it is in the takeover state and one of the nodes in its DR group has the "node-object-limit" MetroCluster configuration option set to off.

### Parameters

(None).

# mcc.switchover events

## mcc.switchover.mdvInNvfailedState

### Severity

ERROR

### Description

This message occurs when a MetroCluster switchover operation finds a system metadata volume, used by config replication applications, in the nvfailed state.

### Corrective Action

Verify the configuration.

### Syslog Message

The 'in-nvfailed-state' was cleared on MetroCluster system metadata volume '%s'. The system will attempt to replicate all recent configuration updates (if any) to the local cluster. After the process completes, verify the configuration.

### Parameters

**Volume** (STRING): Volume

## mcc.toomany events

### mcc.toomany.vols.per.aggr

### Severity

ALERT

### Description

This message occurs when a volume is added to an aggregate in which the recommended maximum number of volumes for an aggregate in a MetroCluster setup is exceeded. An excessive number of volumes in a MetroCluster aggregate can result in high failover times, which may disrupt data availability from the affected node.

### Corrective Action

Re-distribute the volumes among the aggregates to comply with the volume limit per aggregate in a MetroCluster setup. The number of volumes in an aggregate can be found by using the "aggregate show -fields volcount" command. Volumes can be moved between aggregates by using the "volume move start -vserver <vserver\_name> -volume <vol\_name> -destination-aggregate <destination\_aggregate>" command.

### Syslog Message

MetroCluster protected aggregate %s has volume count %d, which exceeds the recommended maximum %d.

### Parameters

**aggregate** (STRING): Name of the aggregate.

**current\_volume** (INT): Current number of volumes on the aggregate.

**volume\_limit** (INT): Recommended maximum number of volumes per aggregate.

# mct events

## mct.channel events

### mct.channel.destroy

**Severity**

NOTICE

**Description**

This message occurs when the Mirror Cache Transport layer destroys a communication channel between two nodes.

**Corrective Action**

(None).

**Syslog Message**

Mirror Cache transport destroyed channel '%s'.

**Parameters**

**channel** (STRING): Name of the data channel.

## mct.control events

### mct.control.error

**Severity**

NOTICE

**Description**

This message occurs when the Mirror Cache Transport layer encounters an error in its control group.

**Corrective Action**

(None).

**Syslog Message**

Mirror Cache transport control error: '%s'.

**Parameters**

**error\_message** (STRING): Details for the encountered error.



# mdb events

## mdb.exit events

### mdb.exit

#### Severity

EMERGENCY

#### Description

This message occurs when the management process exits to recover from a database error.

#### Corrective Action

If the problem persists, contact NetApp technical support for assistance.

#### Syslog Message

MDB %s process exit due to "%s". Error code: %d (%s).

#### Parameters

**name** (STRING): Name of the M-Host process using the management database.

**table\_or\_operation** (STRING): Database table name or operation that encountered the error.

**error\_code** (INT): Numeric error code reported by the management database.

**error\_msg** (STRING): Error message reported by the management database.

# mes\_dimm events

## mes\_dimm.error events

### mes\_dimm.error

#### Severity

ALERT

#### Description

This message appears when an error is encountered by MES (Memory Error Statistics).

#### Corrective Action

DIMM error reported. Please contact NetApp technical support for further assistance.

#### Syslog Message

DIMM error - %s.

#### Parameters

**error** (STRING): Error message from Data ONTAP.

# mgmt events

## mgmt.aba events

### mgmt.aba.abortCloneParent

#### Severity

INFORMATIONAL

#### Description

This message occurs when Auto Balance Aggregate initiates a move of a volume that was then changed to a FlexClone® parent.

#### Corrective Action

(None).

#### Syslog Message

The Auto Balance Aggregate initiated move of volume %s from aggregate %s to aggregate %s was aborted when the move job detected that the volume was now a FlexClone parent.

#### Parameters

**volume** (STRING): Name of the volume.

**srcAggregate** (STRING): Name of the source aggregate.

**dstAggregate** (STRING): Name of the destination aggregate.

### mgmt.aba.abortMirrorMove

#### Severity

INFORMATIONAL

#### Description

This message occurs when Auto Balance Aggregate initiates a move of a volume that was then changed to be part of a SnapMirror® relationship".

#### Corrective Action

(None).

#### Syslog Message

The Auto Balance Aggregate initiated move of volume %s from aggregate %s to aggregate %s was aborted when the move job detected that the volume was now part of a SnapMirror relationship.

#### Parameters

**volume** (STRING): Name of the volume.

**srcAggregate** (STRING): Name of the source aggregate.

**dstAggregate** (STRING): Name of the destination aggregate.

### mgmt.aba.move.unbalanced

## Severity

INFORMATIONAL

## Description

This message occurs when Auto Balance Aggregate is running in auto mode and identifies an aggregate in which the used space has exceeded the high-water mark. Auto Balance Aggregate attempts to trigger moves to lower the aggregate's usage below the threshold.

## Corrective Action

Use the "volume move show" command to monitor active moves.

## Syslog Message

Auto Balance Aggregate has detected that aggregate %s has exceeded the threshold of %jd. The aggregate currently shows a usage of %jd.

## Parameters

**aggr** (STRING): Name of the aggregate.  
**usedSpace** (LONGINT): Aggregate's used space.  
**threshold** (LONGINT): Configured high-water mark.

## mgmt.aba.moveVolume

## Severity

INFORMATIONAL

## Description

This message occurs when Auto Balance Aggregate identifies a volume that can be moved to alleviate a low-space condition with an aggregate. The move will be initiated by Auto Balance Aggregate and can be monitored by using the "volume move show" command.

## Corrective Action

(None).

## Syslog Message

Auto Balance Aggregate is moving volume %s in Vserver %s from aggregate %s to aggregate %s. Upon completion of the move, the source aggregate is projected to be %d percent full and the destination aggregate is projected to be %d percent full.

## Parameters

**volume** (STRING): Volume that will be moved.  
**vserver** (STRING): Vserver hosting the volume.  
**srcAggr** (STRING): Aggregate the volume will be moved from.  
**dstAggr** (STRING): Aggregate the volume will be moved to.  
**afterSourceSpace** (INT): Projected percentage of used space in the source aggregate after the move is performed.  
**afterDestSpace** (INT): Projected percentage of used space in the destination aggregate after the move is performed.

## mgmt.aba.noMovableVols

## Severity

ERROR

## Description

This message occurs when Auto Balance Aggregate identifies an aggregate that is unbalanced but cannot identify a volume to move from that aggregate. Some of the more common reasons this can occur for an aggregate are that its volumes are being heavily accessed, its volumes have been recently moved, or Auto Balance Aggregate cannot identify an aggregate to move to.

## Corrective Action

Manually inspect the aggregate to determine whether any volumes can be moved.

## Syslog Message

Auto Balance Aggregate could not find any volumes to move from aggregate %s. The aggregate's calculated space is %jd and the threshold is %jd.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**usedSize** (LONGINT): Recorded used size.

**threshold** (LONGINT): Configured threshold.

## mgmt.aba.rcmd.unbalanced

## Severity

ERROR

## Description

This message occurs when Auto Balance Aggregate is running in recommend mode and identifies an aggregate in which the used space has exceeded the high-water mark.

## Corrective Action

Use the "volume move recommend show" command to see the list of recommended actions.

## Syslog Message

Auto Balance Aggregate has detected that aggregate %s has exceeded the threshold of %jd. The aggregate currently shows a usage of %jd.

## Parameters

**aggr** (STRING): Name of the aggregate.

**usedSpace** (LONGINT): Aggregate's used space.

**threshold** (LONGINT): Configured high-water mark.

## mgmt.aba.recommendMove

## Severity

INFORMATIONAL

## Description

This message occurs when Auto Balance Aggregate identifies a volume that can be moved to alleviate a low-space condition with an aggregate. The move will not be initiated, but will be listed with other recommendations listed by running the "volume move recommend show" command.

## Corrective Action

If desired, initiate the recommended move by using the "volume move start" command.

## Syslog Message

Auto Balance Aggregate is recommending that volume %s in Vserver %s be moved from aggregate %s to aggregate %s. Upon completion of the move, the source aggregate is projected to be %d percent full and the destination aggregate is projected to be %d percent full.

## Parameters

**volume** (STRING): Volume identified as a candidate to move.

**vserver** (STRING): Vserver hosting the volume.

**srcAggr** (STRING): Aggregate the volume should be moved from.

**dstAggr** (STRING): Aggregate the volume should be moved to.

**afterSourceSpace** (INT): Projected percentage of used space in the source aggregate if the move is performed.

**afterDestSpace** (INT): Projected percentage of used space in the destination aggregate if the move is performed.

# mgmt.abv events

## mgmt.abv.service.filling

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when the percentage of data used on one or more data constituents in a storage service of a managed Infinite Volume exceeds the full watermark. Constituents that have exceeded the full watermark will receive no new content.

### Corrective Action

Either increase the size of this Infinite Volume by entering the "volume modify -size" command, or enter the "autobalance volume start" command to equalize the percentage of data used across all data constituents within this storage service.

### Syslog Message

(None).

### Parameters

**vserver** (STRING): Vserver that hosts the constituent.

**storageService** (STRING): Storage service that holds the constituents.

**fullCount** (INT): Number of constituents that are full.

**totalCount** (INT): Total number of constituents within this storage service.

## mgmt.abv.service.full

## Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

## Severity

ALERT

## Description

This message occurs when the percentage of data used for all of the data constituents within an Infinite Volume's storage service have exceeded the Infinite Volume's full watermark. As a result, users cannot write any more data to this Infinite Volume that would be stored in this storage service by the create-time policy, and the post-ingest policy cannot move files to this storage service while all of its data constituents are in excess of the full watermark.

## Corrective Action

This storage service requires more capacity. Either increase the size of one or more of its data constituents, or add one or more new data constituents to it. After adding more capacity, it is strongly recommended that you enter the "autobalance volume start" command to equalize the percentage of data used across all data constituents in this storage service.

## Syslog Message

The percentage of data used for all of the data constituents within an Infinite Volume's storage service has exceeded the Infinite Volume's full watermark. Vserver: %s, storage service: %s.

## Parameters

**vserver** (STRING): Vserver that hosts these constituents.

**storageService** (STRING): Storage service that holds the constituents.

## mgmt.abv.volume.filling

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when the percentage of data used on one or more data constituents in this Infinite Volume exceeds the full watermark. Constituents that have exceeded the full watermark will receive no new content.

### Corrective Action

Either increase the size of this Infinite Volume by entering the "volume modify -size" command, or enter the "autobalance volume start" command to equalize the percentage of data used across all data constituents.

### Syslog Message

(None).

### Parameters

**vserver** (STRING): Vserver that hosts the constituent.

**fullCount** (INT): Number of full data constituents within this Infinite Volume.

**totalCount** (INT): Total number of data constituents within this Infinite Volume.

## mgmt.abv.volume.full

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ALERT

### Description

This message occurs when all of the data constituents within an Infinite Volume have exceeded the Infinite Volume's full watermark. Users cannot write any more data to this Infinite Volume.

### Corrective Action

Expand this Infinite Volume by entering the "volume modify -size" command. After expanding the Infinite Volume, it is strongly recommended that you enter the "autobalance volume start" command to equalize the percentage of data used across all data constituents.

### Syslog Message

(None).

### Parameters

**vserver** (STRING): Vserver that hosts the constituents.

## mgmt.cifsserver events

### mgmt.cifsserver.failure

#### Severity

ERROR

#### Description

This message occurs when a CIFS server fails to be configured and, therefore, CIFS access might not be available.

#### Corrective Action

Check the reason for the error and take action accordingly. For further assistance, contact technical support.

#### Syslog Message

Failed to configure CIFS server. Error: %s.

#### Parameters

**reason** (STRING): Reason for the failure.

**vservername** (STRING): Name of the vservers.

**cifsservername** (STRING): Name of the cifs server.

### mgmt.cifsserver.success

#### Severity

INFORMATIONAL



## Description

This message occurs when a CIFS server is successfully configured after a failure.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**vservername** (STRING): Name of the vserver.

**cifsservername** (STRING): Name of the cifs server.

# mgmt.cifsshare events

## mgmt.cifsshare.duplicatesid

### Severity

NOTICE

### Description

This message occurs when a network client attempts to set a CIFS share's discretionary access control list (DACL) to a list where any security descriptor (SID) appears in more than one access control entry (ACE). The operation is failed and no changes are made to the share's DACL.

### Corrective Action

(None).

### Syslog Message

Duplicate SID not supported in CIFS share DACL %s

### Parameters

**vservername** (STRING): Name of the Vserver.

**cifssharename** (STRING): Name of the CIFS share.

**sid** (STRING): Textual SID.

## mgmt.cifsshare.failure

### Severity

ERROR

### Description

This message occurs when configuration of a CIFS share fails; CIFS access to the share might not be available.

### Corrective Action

Check the reason for the error and take action accordingly.

### Syslog Message

Failed to configure CIFS share. Error: %s.

## Parameters

**reason** (STRING): Reason for the failure.  
**vservername** (STRING): Name of the vserver.  
**cifsservername** (STRING): Name of the cifs server.  
**cifssharename** (STRING): Name of the cifs share.

## mgmt.cifsshare.success

### Severity

INFORMATIONAL

### Description

This message occurs when a CIFS share is successfully configured after a failure.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vservername** (STRING): Name of the vserver.  
**cifssharename** (STRING): Name of the cifs share.

## mgmt.cifssuperuser events

### mgmt.cifssuperuser.failure

### Severity

ERROR

### Description

This message occurs when configuration of a CIFS super-user fails; CIFS access through this account may not grant full super-user privileges.

### Corrective Action

Check the reason for the error and take action accordingly.

### Syslog Message

Failed to configure CIFS superuser. Error: %s.

## Parameters

**reason** (STRING): Reason for the failure.  
**vservername** (STRING): Name of the vserver.  
**cifsservername** (STRING): Name of the cifs server.  
**cifsdomainname** (STRING): Name of the cifs domain.  
**accountname** (STRING): Name of the superuser account.

## mgmt.cifssuperuser.success

### Severity

INFORMATIONAL

### Description

This message occurs when a CIFS super-user is successfully configured after a failure.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**vservername** (STRING): Name of the vserver.  
**cifsservername** (STRING): Name of the cifs server.  
**cifsdomainname** (STRING): Name of the cifs domain.  
**accountname** (STRING): Name of the superuser account.

## mgmt.cifssymlink events

### mgmt.cifssymlink.failure

#### Severity

ERROR

#### Description

This message occurs when configuration of a CIFS symbolic link fails; the symbolic link might not be available through CIFS.

#### Corrective Action

Check the reason for the error and take action accordingly.

#### Syslog Message

Failed to configure CIFS symlink. Error: %s.

#### Parameters

**reason** (STRING): Reason for the failure.  
**vservername** (STRING): Name of the vserver.  
**unix\_path** (STRING): The UNIX path.

### mgmt.cifssymlink.success

#### Severity

INFORMATIONAL

#### Description

This message occurs when a CIFS symbolic link is successfully configured after a failure.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**vservername** (STRING): Name of the vserver.

**unix\_path** (STRING): The UNIX path.

## mgmt.cr events

### mgmt.cr.constituent.added

**Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

INFORMATIONAL

**Description**

This message occurs when a constituent volume is added to the Infinite Volume.

**Corrective Action**

(None).

**Syslog Message**

Constituent volume "%s" has been added to Infinite Volume "%s" in Vserver "%s". Its size is %ld bytes.

**Parameters**

**constituentName** (STRING): Name of the added volume.

**repositoryName** (STRING): Name of the Infinite Volume.

**vserverName** (STRING): Name of the Vserver.

**size** (LONGINT): Size of the added volume.

### mgmt.cr.constituent.deleted

**Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

**Severity**

INFORMATIONAL

**Description**

This message occurs when a constituent volume is deleted from the Infinite Volume.

**Corrective Action**

(None).

## Syslog Message

Constituent volume "%s" has been deleted from the Infinite Volume "%s" in Vserver "%s".

## Parameters

**constituentName** (STRING): Name of the deleted volume.

**repositoryName** (STRING): Name of the Infinite Volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.cr.constituent.failure

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when an Infinite Volume operation that operates on the constituent volumes succeeds, but fails on an individual constituent.

### Corrective Action

Fix the cause of the failure listed in the error, and then run the command again.

## Syslog Message

Repository job "%s" run on Infinite Volume "%s" in Vserver "%s" failed for constituent volume "%s": %s

## Parameters

**jobName** (STRING): Name of the repository job.

**repositoryName** (STRING): Name of the Infinite Volume.

**vserverName** (STRING): Name of the Vserver.

**constituentName** (STRING): Name of the constituent volume.

**reason** (STRING): Reason for the failure.

## mgmt.cr.constituent.resized

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when an Infinite Volume constituent volume is resized.

### Corrective Action

(None).

## Syslog Message

Constituent volume "%s" of Infinite Volume "%s" in Vserver "%s" has been resized to %ld bytes.

## Parameters

**constituentName** (STRING): Name of the resized volume.  
**repositoryName** (STRING): Name of the Infinite Volume.  
**vserverName** (STRING): Name of the Vserver.  
**size** (LONGINT): New size of the volume.

## mgmt.cr.job.failed

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when a repository job fails.

### Corrective Action

Fix the cause of the failure, listed in the error, and run the command again.

### Syslog Message

Repository job "%s" has failed for Infinite Volume "%s" in Vserver "%s": %s

## Parameters

**jobName** (STRING): Name of the repository job.  
**repositoryName** (STRING): Name of the Infinite Volume.  
**vserverName** (STRING): Name of the Vserver.  
**reason** (STRING): Reason for the failure.

## mgmt.cr.job.succeeded

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when a repository job is completed successfully.

### Corrective Action

(None).

### Syslog Message

Repository job "%s" has succeeded for Infinite Volume "%s" in Vserver "%s": %s

## Parameters

**jobName** (STRING): Name of the repository job.  
**repositoryName** (STRING): Name of the Infinite Volume.  
**vserverName** (STRING): Name of the Vserver.  
**reason** (STRING): Success message of the job.

## mgmt.cr.missing.footprint

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when Data ONTAP® is attempting to create or expand an Infinite Volume and cannot retrieve volume footprint information about a volume. The footprint information for all volumes on aggregates used by an Infinite Volume, even those in other Vservers, must be known to safely create or expand an Infinite Volume.

### Corrective Action

Ensure that the volume is healthy and online, and then retry the operation.

### Syslog Message

Unable to get footprint information about volume "%s" in Vserver "%s".;

### Parameters

**constituentName** (STRING): Name of the volume for which footprint information could not be retrieved.  
**vserverName** (STRING): Name of the Vserver the volume belongs to.

## mgmt.exovol events

### mgmt.exovol.transfer.failed

### Severity

ERROR

### Description

This message occurs when the replication engine fails to transfer exo-volume data.

### Corrective Action

If problem persists, Contact NetApp technical support.

### Syslog Message

The replication for volume dsid %jd, transfer-id %s failed to transfer %s exo-volume data, error: %d, failure reason: %s, extended failure reason: %s.

### Parameters

**volDsid** (LONGINT): Data set identifier of the volume.  
**transferId** (STRING): Replication transfer id.  
**streamType** (STRING): Type of data being transferred.  
**error** (INT): Transfer failure error code.  
**failureReason** (STRING): Basic reason for the failure.  
**extendendFailureReason** (STRING): Detailed reason for the failure.

## mgmt.exovol.xfer.dest.no.space

### Severity

ERROR

### Description

This message occurs when there is no free space left on the destination aggregate.

### Corrective Action

Please increase the amount of free space available on the destination aggregate.

### Syslog Message

The replication for volume dsid %jd, transfer-id %s failed to transfer %s exo-volume data because there is no free space left on the destination aggregate.

### Parameters

**volDsid** (LONGINT): Data set identifier of the volume.

**transferId** (STRING): Replication transfer id.

**streamType** (STRING): Type of data being transferred.

## mgmt.fpolicy events

### mgmt.fpolicy.policy.disabled

### Severity

INFORMATIONAL

### Description

This message occurs when a Vserver administrator disables an FPolicy® policy. No subsequent client requests will be processed and no notifications will be sent to FPolicy servers attached with this policy.

### Corrective Action

(None).

### Syslog Message

FPolicy policy %s is disabled on Vserver %s.

### Parameters

**policyname** (STRING): Name of the policy.

**vservername** (STRING): Name of the Vserver.

### mgmt.fpolicy.policy.enabled

### Severity

INFORMATIONAL

### Description

This message occurs when a Vserver administrator enables an FPolicy® policy. Based on the event configuration, subsequent client requests will be processed, notifications will be generated, and notifications might be sent (based on the policy configuration) to FPolicy servers attached with this policy.



**Corrective Action**

(None).

**Syslog Message**

FPolicy policy %s is enabled on Vserver %s.

**Parameters**

**polycname** (STRING): Name of the policy.

**vservername** (STRING): Name of the Vserver.

**mgmt.fpolicy.priv.failed****Severity**

ERROR

**Description**

This message occurs when the creation of privileged share ONTAP\_ADMIN\$ fails. Attempted connections to ONTAP\_ADMIN\$ share by the FPolicy servers will fail.

**Corrective Action**

Check the reason for failure and take action accordingly. If the problem persists, contact NetApp technical support for assistance.

**Syslog Message**

Failed to create privileged share ONTAP\_ADMIN\$ for Vserver "%s". Reason: "%s".

**Parameters**

**vserverName** (STRING): Name of the Vserver.

**reason** (STRING): Reason for failure.

**mgmt.fpolicy.replay.failed****Severity**

ERROR

**Description**

This message occurs when an FPolicy® configuration replication mechanism fails after a maximum number of retries.

**Corrective Action**

Restart the node or contact NetApp technical support for assistance.

**Syslog Message**

FPolicy configuration replication process failed.

**Parameters**

(None).

# mgmt.fsecurity events

## mgmt.fsecurity.failure

### Severity

ERROR

### Description

This message occurs when applying configuration of fsecurity fails.

### Corrective Action

Wait for a while and reapply using the "vserver security file-directory apply" command. See the "vserver security file-directory" man pages for information, or if the problem persists, contact NetApp technical support.

### Syslog Message

Failed to apply file security policy. Error: %s.

### Parameters

**reason** (STRING): Reason for the failure.

**vservername** (STRING): Name of the vserver.

## mgmt.fsecurity.success

### Severity

INFORMATIONAL

### Description

This message occurs when a file security policy was applied successfully after a previous failure.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**vservername** (STRING): Name of the Vserver.

# mgmt.keymgr events

## mgmt.keymgr.keyid.noDelete

### Severity

ERROR

### Description

This message occurs when the system cannot delete the key ID associated with the encrypted aggregate from the key management server.

### Corrective Action

Manually delete the key ID associated with the encrypted aggregate from the key management server.

### Syslog Message

Encrypted aggregate "%s" with UUID "%s" has been deleted. The key ID "%s" was associated with this aggregate and must be manually deleted from the key management server.

### Parameters

**name** (STRING): Aggregate name.

**uuid** (STRING): Unique identifier of the aggregate when the aggregate name by itself is insufficient.

**key\_id** (STRING): Unique key ID used by the encrypted aggregate.

## mgmt.mcc events

### mgmt.mcc.fsecurity.config.err

#### Severity

ERROR

#### Description

This message occurs when MetroCluster(tm) switchback or switchover happens while applying a file security (fsecurity) policy on the MetroCluster partner. The policy was not applied to the partner.

#### Corrective Action

Reapply the policy by using the "vserver security file-directory apply" command. See the "vserver security file-directory" man pages for information. If the problem persists, contact NetApp technical support.

#### Syslog Message

Failed to apply file security policy "%s" on Vserver "%s".

#### Parameters

**policy** (STRING): Fsecurity policy name.

**vservername** (STRING): Name of the Vserver.

## mgmt.nfslisteners events

### mgmt.nfslisteners.failure

#### Severity

ERROR

#### Description

This message occurs when the listener for the NFS service is not configured. If the problem persists, NFS access is not available.

#### Corrective Action

This problem might be resolved by subsequent automatic retries of the NFS request. If NFS service is not available, check the NFS configuration of the service and the external services it uses. To verify NFS configuration of the service, use the "nfs server show -vserver <vserver name> [-fields <field name>]" command. If the problem persists after you check the NFS configuration, delete the NFS Vserver, and then

re-create it. See the NFS man page guide for information about deleting and re-creating the NFS Vserver, or contact NetApp technical support.

### **Syslog Message**

Failed to configure NFS listener. Error: %s.

### **Parameters**

**reason** (STRING): Message provided by the NFS service when a listener configuration does not succeed.

## **mgmt.nfslisteners.success**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the listener for the NFS service is configured after a failure.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

(None).

## **mgmt.nfsserver events**

### **mgmt.nfsserver.failure**

### **Severity**

ERROR

### **Description**

This message occurs when a NFS server fails to be configured and, therefore, NFS access might not be available.

### **Corrective Action**

Check the reason for the error and take action accordingly. For further assistance, contact technical support.

### **Syslog Message**

Failed to configure NFS server. Error: %s.

### **Parameters**

**reason** (STRING): Reason for the failure.

**vservername** (STRING): Name of the vservers.

### **mgmt.nfsserver.success**

**Severity**

INFORMATIONAL

**Description**

This message occurs when a NFS server is successfully configured after a failure.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**vservname** (STRING): Name of the nfs server.

## mgmt.rst events

### mgmt.rst.fail.dst.compressed

**Severity**

ERROR

**Description**

This message occurs when the SnapMirror® 'restore' command fails because compression is enabled on the destination volume.

**Corrective Action**

Force a cleanup of the failed operation by reissuing the 'restore' command with the parameter 'clean-up-failure', disable compression on the destination volume, create a Snapshot(tm) copy, and then reissue the 'restore' command.

**Syslog Message**

SnapMirror restore from source volume '%s' to destination volume '%s' failed with error: %s.

**Parameters**

**srcpath** (STRING): Path name of the source volume for restore.

**dstpath** (STRING): Path name of the destination volume for restore.

**error** (STRING): Error message.

## mgmt.s3 events

### mgmt.s3.sm.license.near.limit

**Severity**

NOTICE

**Description**

This message occurs when the total cluster-wide logical used size for S3 Snapmirror external object store relationships is near the license capacity limit but does not exceed the limit.

### **Corrective Action**

To increase the license capacity, install a new license on the cluster. Alternatively, remove some data you are managing.

### **Syslog Message**

%s.

### **Parameters**

**message** (STRING): Capacity warning message for users

## **mgmt.s3.sm.license.over.limit**

### **Severity**

ERROR

### **Description**

This message occurs when the total cluster-wide bucket used size for S3 Snapmirror external object store relationships is over the license capacity limit.

### **Corrective Action**

To increase the license capacity, install a new license on the cluster. Alternatively, remove some data you are managing.

### **Syslog Message**

%s.

### **Parameters**

**message** (STRING): Capacity warning message for users

## **mgmt.sm events**

### **mgmt.sm.cant.clear.autoinit**

### **Severity**

NOTICE

### **Description**

This message occurs when a SnapMirror® protect operation fails because the source volume is already protected with a SnapMirror relationship in the specified destination Vserver with the same policy, schedule, relationship type, destination volume prefix, and destination volume suffix and auto-initialize parameter set to true. The auto-initialize parameter cannot be cleared on a SnapMirror relationship once it is set to true.

### **Corrective Action**

(None).

### **Syslog Message**

Job ID %lld failed to disable auto-initialize for SnapMirror relationship on volume %s, destination vserver %s. This operation is not permitted.

## Parameters

**jobID** (LONGINT): Job ID.

**srcpath** (STRING): Path name of the source volume to be protected.

**dstvserver** (STRING): Destination Vserver name.

## mgmt.sm.cg.force.delete

### Severity

ERROR

### Description

This message occurs when a sync granular consistency group relationship is being deleted with the force option. Consistency Group relationship will be the deleted, however the items of the Consistency Group might not be made read-write or deletable after the operation.

### Corrective Action

Perform a volume offline and online on the destination volume to trigger the cleanup of any remaining relationship information on the volume.

### Syslog Message

Sync granular consistency group relationship with source path %s and destination path %s has been force deleted and items of the consistency group on the destination volume might not be read-write or deletable.

## Parameters

**src\_path** (STRING): Source path of the sync granular consistency group relationship that was force deleted.

**dst\_path** (STRING): Destination path of the sync granular consistency group relationship that was force deleted.

## mgmt.sm.dp.depr.msg

### Severity

NOTICE

### Description

SnapMirror relationships of type DP are deprecated. From the previous release, users were no longer able to initialize new relationships of type DP. From this release, relationships of type DP will no longer function.

### Corrective Action

Use SnapMirror relationships of type XDP instead of type DP.

### Syslog Message

SnapMirror relationships of type DP are deprecated in this ONTAP release.

## Parameters

(None).

## mgmt.sm.ep.cant.clr.autolnit

### Severity

NOTICE

## Description

This message occurs when a SnapMirror® protect operation fails because the source is already protected with a SnapMirror relationship that has the specified destination endpoint type and the same "-destination-volume-prefix" and "-destination-volume-suffix" parameter values, with the "-auto-initialize" parameter set to "true". The "-auto-initialize" parameter cannot be cleared on a SnapMirror relationship after it is set to "true".

## Corrective Action

To change the auto-initialize behavior, delete and release the existing SnapMirror relationship by using the "snapmirror delete" command and the "snapmirror release" command, and then establish another relationship without auto-initialize.

## Syslog Message

Job ID %lld: Failed to disable auto-initialize for SnapMirror relationship from source %s to destination endpoint type %s. This operation is not permitted.

## Parameters

**jobID** (LONGINT): Job ID.  
**srcpath** (STRING): Path name of the source to be protected.  
**dsteptype** (STRING): Destination endpoint type.

## mgmt.sm.expand.overprovision

### Severity

ERROR

### Description

This message occurs when a volume expand job starts for the source of a SnapMirror® relationship and the provisioning check finds that there isn't enough free space on the aggregates for destination volume. The source volume expand job is not impacted, but auto-expand will fail for the destination FlexGroup after the source FlexGroup is expanded.

### Corrective Action

Increase the size of the aggregates for the destination Flexgroup volume.

### Syslog Message

The aggregates on the cluster hosting destination FlexGroup volume '%s' do not have enough free space to be used for expanding the FlexGroup volume. Expansion of the destination will fail until the size of enough aggregates is increased to provide enough space.

### Parameters

**dstVolume** (STRING): Destination volume name of the SnapMirror® relationship.

## mgmt.sm.keymgr.not.setup

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation on an encrypted source volume fails because the Key Manager is not set up on the destination node. A SnapMirror relationship with the destination in the specified destination Vserver to protect the source volume will not be created until this condition is corrected.



### Corrective Action

Set up the Key Manager by using the "security key-manager setup" command on the destination.

### Syslog Message

SnapMirror protect of source volume "%s" failed to destination Vserver "%s" because the Key Manager was not set up on the destination cluster. Job ID: %llu.

### Parameters

**srcpath** (STRING): Path name of the source volume to be protected.

**dstvserver** (STRING): Destination Vserver name.

**jobID** (LONGINT): Job ID.

## mgmt.sm.license.near.limit

### Severity

NOTICE

### Description

This message occurs when the total cluster-wide physical used size for object store relationships is near the license capacity limit but does not exceed the limit.

### Corrective Action

To increase the license capacity, install a new license on the cluster. Alternatively, remove some data you are managing.

### Syslog Message

%s.

### Parameters

**message** (STRING): Capacity warning message for users

## mgmt.sm.license.not.found

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation of a source volume fails because a required license on the destination cluster is missing. The error message contains the information about the missing license. A SnapMirror relationship with the destination in the specified destination Vserver to protect the source volume will not be created until this condition is corrected.

### Corrective Action

Install the license that is missing on the destination cluster. The error message contains information about the missing license.

### Syslog Message

SnapMirror protect of source volume "%s" failed to destination Vserver "%s" because there was a license missing on the destination cluster. Error: %s, Job ID: %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume to be protected.  
**dstvserver** (STRING): Destination Vserver name.  
**error** (STRING): Error message.  
**jobID** (LONGINT): Job ID.

## mgmt.sm.license.over.limit

### Severity

ERROR

### Description

This message occurs when the total cluster-wide physical used size for object store relationships is over the license capacity limit.

### Corrective Action

To increase the license capacity, install a new license on the cluster. Alternatively, remove some data you are managing.

### Syslog Message

%s.

## Parameters

**message** (STRING): Capacity warning message for users

## mgmt.sm.obj.fanout.unsupp

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails because the source is already protected with a SnapMirror relationship with an object store destination. Creation of multiple SnapMirror relationships from a source endpoint to object store destinations is not supported.

### Corrective Action

View the existing SnapMirror relationship using the "snapmirror show" command, and correct as necessary.

### Syslog Message

Source endpoint "%s" is already protected in an object store. Job ID: %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume to be protected.  
**jobID** (LONGINT): Job ID.

## mgmt.sm.protect.internal.err

### Severity

ERROR

## Description

This message occurs when a SnapMirror® protect operation fails because of an unexpected internal error.

## Corrective Action

Retry the command after some time. If the problem persists, contact NetApp technical support.

## Syslog Message

SnapMirror protect failed because of an unexpected internal error. Error: %s, Job ID: %llu.

## Parameters

**error** (STRING): Error message.

**jobID** (LONGINT): Job ID.

## mgmt.sm.rel.exists

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails because the source volume is already protected with a SnapMirror relationship in the specified destination Vserver with the same policy, schedule, relationship type, destination volume prefix, and destination volume suffix.

### Corrective Action

If you want to protect the volume again, 1. Use a different destination volume suffix or destination volume prefix. 2. Use a different schedule or policy. 3. Use the "volume create" command to create a destination volume with a user defined name and then the "snapmirror create" command to create another SnapMirror relationship for the same source volume.

### Syslog Message

SnapMirror protect operation of source volume "%s" failed to destination Vserver "%s" because the source volume is already protected with a SnapMirror relationship in the specified destination Vserver with the same policy, schedule, relationship type, destination volume prefix, and destination volume suffix. Error: %s, Job ID: %llu.

### Parameters

**srcpath** (STRING): Path name of the source volume to be protected.

**dstvserver** (STRING): Destination Vserver name.

**error** (STRING): Error message.

**jobID** (LONGINT): Job ID.

## mgmt.sm.rename.ndu.error

### Severity

ERROR

### Description

This message occurs when the renaming of a SnapMirror® volume in another cluster fails, due to the lack of support for this operation in that cluster's effective version of ONTAP® software.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Failed to update Volume name "%s" in remote cluster "%s".

**Parameters**

**volume** (STRING): Volume for which the rename operation failed.

**clusName** (STRING): Remote cluster name.

**mgmt.sm.restore.fail****Severity**

ERROR

**Description**

This message occurs when the SnapMirror® 'restore' command fails.

**Corrective Action**

To retry the restore, reissue the command. Alternatively, force a cleanup of the failed operation without starting another restore by including the parameter 'clean-up-failure' when reissuing the 'restore' command.

**Syslog Message**

SnapMirror restore from source volume '%s' to destination volume '%s' failed with error: %s.

**Parameters**

**srcpath** (STRING): Path name of the source volume for restore.

**dstpath** (STRING): Path name of the destination volume for restore.

**error** (STRING): Error message.

**mgmt.sm.schd.invalid****Severity**

ERROR

**Description**

This message occurs when the SnapMirror® protect operation fails to find the specified schedule name or the schedule has not been specified. A SnapMirror relationship with the destination will be created and initialized, if auto-initialize is enabled, but subsequent periodic backup operations will not be initiated.

**Corrective Action**

Add a valid schedule.

**Syslog Message**

The schedule "%s" is not valid or has not been specified.

**Parameters**

**schd\_name** (STRING): Name of the schedule that is invalid.

## mgmt.sm.schd.unavailable

### Severity

ERROR

### Description

This message occurs when a schedule is not available on the secondary site for configuration replication of a SnapMirror® relationship of -type Load Sharing (LS) in a MetroCluster(tm) configuration.

### Corrective Action

Create a job schedule on the secondary site that is identical to that on the primary site by using the 'job schedule cron create' command.

### Syslog Message

Configuration replication of SnapMirror relationship of -type Load Sharing for destination path "%s" has failed because the schedule "%s" is not available.

### Parameters

**dstpath** (STRING): Destination path name of the SnapMirror relationship.

**schd\_name** (STRING): Name of the schedule that is unavailable on the secondary site.

## mgmt.sm.snap.create.fail

### Severity

ALERT

### Description

This message occurs when a SnapMirror® "snapshot create" task fails on the destination volume, after successfully converting a FlexVol® volume to a FlexGroup. Future SnapMirror transfers will be unsuccessful.

### Corrective Action

Create the baseline Snapshot copy on destination volume "%s" using the "volume conversion create-destination-snapshot" command, and then re-create the SnapMirror relationship using the "snapmirror create -destination-path %s -source-path %s -policy %s -throttle %d -schedule %s" command.

### Syslog Message

Create the baseline Snapshot copy on destination volume "%s" using the "volume conversion create-destination-snapshot" command, and then re-create the SnapMirror relationship using the "snapmirror create -destination-path %s -source-path %s -policy %s -throttle %s -schedule %s" command.

### Parameters

**volume** (STRING): Volume converted to a FlexGroup.

**dstPath** (STRING): Destination path for SnapMirror create.

**srcPath** (STRING): Source path for SnapMirror create.

**policy** (STRING): Policy for SnapMirror create.

**throttle** (STRING): Throttle for SnapMirror create.

**schedule** (STRING): Schedule for SnapMirror create.

## mgmt.sm.src.policy.creat

## Severity

ALERT

## Description

This message occurs when a reverse resync operation was performed from System Manager and the SnapMirror policy associated with an intercluster SnapMirror relationship was replicated to the source from the destination, but the retention period in the policy was not replicated. This will happen only if the source cluster ONTAP® version is earlier than ONTAP 9.13.1. Because the retention period was not associated with the Snapshot copies, the Snapshot copies are prone to deletion, which might lead to data protection issues.

## Corrective Action

Add the retention period to the appropriate rules of the SnapMirror policy in the source cluster using the command "snapmirror policy modify-rule -vserver <storage\_VM\_name> -policy <policy\_name> -snapmirror -label <snapmirror\_label\_name> -retention-period <retention\_period>".

## Syslog Message

The retention period setting in the destination SnapMirror policy "%s" has not been replicated to the source during a reverse resync, and it might need to be updated manually.

## Parameters

**policyName** (STRING): Name of the policy associated with the SnapMirror relationship that was replicated from the destination to source.

## mgmt.sm.src.vol.not.found

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails because the source volume was not found.

### Corrective Action

Check whether the source path is correct, and then retry the operation with the correct path.

### Syslog Message

SnapMirror protect of source volume "%s" failed to destination Vserver "%s" because the source volume was not found. Job ID: %llu.

### Parameters

**srcpath** (STRING): Path name of the source volume to be protected.

**dstvserver** (STRING): Destination Vserver name.

**jobID** (LONGINT): Job ID.

## mgmt.sm.sync.break.dp.error

### Severity

ERROR

## Description

This message occurs when a SnapMirror® break operation fails on a synchronous SnapMirror destination volume, because the destination is either in the "broken-off" state, or another state that is inconsistent with the operation.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

SnapMirror break operation failed for sync SnapMirror relationship with destination volume "%s" because the destination volume is either in the broken-off state or the volume state is inconsistent.

## Parameters

**dstpath** (STRING): Path name of the destination volume.

## mgmt.sm.v2.upgrade.chk

### Severity

ERROR

### Description

This message occurs after booting, when there is a Data Protection relationship created in a Data ONTAP® release prior to 8.2.

### Corrective Action

Establish Vserver peering so that the system can convert the pre-8.2 relationships to 8.2 and later relationships. Vserver peering can only be established if both endpoints are running Data ONTAP 8.2 or later on all nodes in the cluster. To check whether the source cluster supports Data ONTAP 8.2 SnapMirror® relationships, run the "capability show -node \* -capability volume.snapmirror\_v2" command. If the capability is "true" on all nodes, create a Vserver peer relationship between the source and destination endpoint Vservers by issuing the "vserver peer create" command. If the capability is "false" on any node in the source cluster, upgrade and configure the source cluster by following instructions in the Upgrade and Revert/Downgrade Guide.

### Syslog Message

SnapMirror relationships from a pre-8.2 installation are present on this cluster but no Vserver peer relationships were found between the source and destination Vservers.

### Parameters

(None).

## mgmt.sm.vol.place.create.err

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails to find an aggregate with sufficient free space for a destination volume or fails to create the volume. A SnapMirror relationship with the destination in the specified destination Vserver to protect the source volume will not be created until this condition is corrected.

## Corrective Action

Use the "vserver show -vserver <vserver-name> -fields aggr-list" command to check whether there is a list of aggregates assigned to the destination Vserver. If there are aggregates listed, check the status of and the space left in the storage aggregates belonging to that Vserver. If there are no aggregates listed, then check the status of and the space left in all of the non-root, non-taken-over and non-SnapLock aggregates. To check whether storage aggregates are non-root, non-taken-over and non-SnapLock, run the "storage aggregate show -fields root, is-home, snaplock-type" command. To check the status of and the space left in a storage aggregate, run the "storage aggregate show -aggregate <aggr-name>" command for a particular aggregate or "storage aggregate show" for all storage aggregates on the cluster. If storage aggregates are offline, use "storage aggregate online -aggregate <aggr-name>" command to bring the storage aggregates online. If there is no space in the storage aggregates: 1) Use the "vserver add-aggregates -vserver <vserver-name> -aggregates <aggr-name>" command to add more storage aggregates to the destination Vserver. 2) If needed, use the "storage aggregate create -aggregate <aggr-name> -diskcount <number-of-disks> -node <node-name>" command to create new storage aggregates on available nodes. 3) Use the "storage aggregate add-disks -aggregate <aggr-name> -diskcount <number-of-disks>" command to add space to the storage aggregate. Even if there is free space on storage aggregates, it is possible that the maximum number of volumes already exist on a node. If so, additional volumes cannot be created on that node. If this is the case, add more aggregate space on other nodes. You can determine the number of volumes on a node by using the "volume show -node <node-name>" command. Note the entry count at the end of the show output. To check the overall status of the SnapMirror protect operation, use the "job show <jobID> -instance" or "job history show -id <jobID> -instance" commands. If the problem persists, contact NetApp technical support.

## Syslog Message

SnapMirror protect of source volume "%s" failed on destination Vserver "%s". Error: %s, Job ID: %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume to be protected.  
**dstvserver** (STRING): Destination Vserver name.  
**error** (STRING): Error message.  
**jobID** (LONGINT): Job ID.

## mgmt.sm.vol.style.unsupp

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails because the source volume style is not supported. The SnapMirror protect command only supports FlexVol® volumes. FlexGroups, FlexGroup constituents, Infinite Volumes, and Infinite Volume constituents are not supported. Temporary volumes are also not supported.

## Corrective Action

If you want to protect Infinite Volumes or FlexGroups, you must use the individual commands: 1. "volume create" of the destination volume 2. "snapmirror create" of the relationship 3. "snapmirror initialize" of the relationship Creating individual SnapMirror relationships on FlexGroup constituents or Infinite Volume constituents is not supported.

## Syslog Message

SnapMirror protect of source volume "%s" failed because the volume is not a FlexVol volume. Job ID: %llu.



## Parameters

**srcpath** (STRING): Path name of the source volume to be protected.

**jobID** (LONGINT): Job ID.

## mgmt.sm.vsrs.not.peered

### Severity

ERROR

### Description

This message occurs when a SnapMirror® protect operation fails because either the source or destination Vservers do not exist or the source and destination Vservers are not peered. A SnapMirror relationship with the specified destination Vserver will not be created, and the source volume will not be protected, until this condition is corrected.

### Corrective Action

Use the 'vserver show' and 'vserver peer' commands to verify that the source and destination Vservers exist and are named uniquely. Also make sure that the 'snapmirror protect' operation targets the intended Vservers. See the man pages of these commands for further information.

### Syslog Message

SnapMirror protect operation of source volume "%s" failed because destination Vserver "%s" and source Vserver %s are not peered. Job ID: %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume to be protected.

**srcvserver** (STRING): Source Vserver name.

**dstvserver** (STRING): Destination Vserver name.

**jobID** (LONGINT): Job ID.

## mgmt.smbc events

### mgmt.smbc.config.not.repl

### Severity

NOTICE

### Description

This message occurs when the system fails to replicate the configuration of Report Target Group Port(RTPG) and Inter-Cluster(IC) LIF information.

### Corrective Action

(None).

### Syslog Message

Failure in configuration replication of RTPG, IC LIF information with relationship UUID "%s".

## Parameters

**srcpath** (STRING): Path name of the source group of the SnapMirror Business Continuity relationship.

**dstpath** (STRING): Path name of the destination group of the SnapMirror Business Continuity relationship.

**relationship\_id** (STRING): Relationship UUID.

## mgmt.smbc.proxy.create.fail

### Severity

ERROR

### Description

This message occurs when there is a failure to create proxy I/O path channel between the SMBC peer clusters.

### Corrective Action

Check the network connection between the source and destination volumes by using the "net interface show -role intercluster -status-admin up" command.

### Syslog Message

Failure in proxy initiator %s and target %s creation.

### Parameters

**proxy\_initiator** (STRING): Path name of the proxy initiator.

**proxy\_target** (STRING): Path name of the proxy target.

**relationship\_id** (STRING): Relationship UUID.

**error\_msg** (STRING): Description of the failure.

## mgmt.snapmir events

### mgmt.snapmir.dst.invalidType

#### Severity

ERROR

#### Description

This message occurs when the destination volume of a SnapMirror® relationship is not of -type Data Protection (DP) or Load Sharing (LS), and a 'snapmirror initialize' operation is attempted. A volume must be of -type DP or LS to accept a 'snapmirror initialize' operation.

#### Corrective Action

Use a different volume of -type DP or LS, or use the 'volume create' command to create a new volume of -type DP or LS to use as a destination of the SnapMirror relationship. You cannot modify the type of a volume using the 'volume modify' command.

#### Syslog Message

Destination volume '%s' of a SnapMirror relationship is not of -type DP or LS. Job ID %llu.

#### Parameters

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**jobID** (LONGINT): Job ID.

### mgmt.snapmir.fl.bkt.del.fail

#### Severity

ERROR

## Description

This message occurs when a bucket could not be deleted during a cleanup of a SnapMirror restore or create operation. Because the operation did not complete successfully, the bucket might be left in an inconsistent state.

## Corrective Action

Use the "object-store-server bucket delete" command to delete the bucket.

## Syslog Message

Bucket "%s" could not be deleted as part of the clean up of a SnapMirror create or restore operation.

## Parameters

**bucket** (STRING): Name of the bucket in the SnapMirror relationship.

## mgmt.snapmir.fl.crt.cleanup

### Severity

ERROR

### Description

This message occurs when the cleanup of a link failed when it encountered an error during a SnapMirror create operation. Because the link could not be cleaned up properly, the bucket could be in an inconsistent state and might not be able to accept any traffic. Also, if the links are not cleaned up, a new SnapMirror relationship cannot be established with these two buckets.

### Corrective Action

1. Run the "snapmirror delete" command to delete the link.

### Syslog Message

Cleanup of the link between buckets "%s" and "%s" failed with error "%s".

### Parameters

**sourceBucket** (STRING): Name of the source bucket in the SnapMirror relationship.

**destBucket** (STRING): Name of the destination bucket in the SnapMirror relationship.

**error** (STRING): Error that was encountered.

## mgmt.snapmir.fl.failover.rmt

### Severity

ERROR

### Description

This message occurs when the remote bucket was not demoted during a failover operation. This could allow the remote bucket to still accept the traffic even though it is no longer permitted to so. Use the "demote bucket" command to avoid this situation.

### Corrective Action

Use the "fabriclink demote-bucket" command in the cluster where remote bucket "%s" is located.

### Syslog Message

Failover failed for remote bucket "%s".

## Parameters

**remoteBucket** (STRING): Name of the remote bucket for which the failover operation failed.

## mgmt.snapmir.fl.mod.rmt.fail

### Severity

NOTICE

### Description

This message occurs when the links on a remote cluster could not be modified during the execution of the "snapmirror modify" command. Even though the SnapMirror policy was modified, some of the links on the remote cluster could not be updated with new rpo/ throttle values. There is no detectable impact until a failover is performed. Only when a failover is performed would the reverse link be used, and it would still be functional even if the links on the remote cluster were not updated previously. They would just be using old rpo/throttle values.

### Corrective Action

(None).

### Syslog Message

Links on remote cluster "%s" were not modified.

## Parameters

**clusterUuid** (STRING): Cluster UUID of the cluster where the links were not updated.

## mgmt.snapmir.fl.rstr.bkt.cln

### Severity

ERROR

### Description

This message occurs when objects in a restore bucket were not deleted during the cleanup of a SnapMirror restore operation. The restore bucket has objects in it even though the restore relationship no longer exists.

### Corrective Action

Delete the objects in bucket "%s".

### Syslog Message

Objects in bucket "%s" were not deleted during a restore failure cleanup.

## Parameters

**restoreBucket** (STRING): Name of the restore bucket in the SnapMirror relationship.

## mgmt.snapmir.fl.rstr.cln.all

### Severity

ERROR

### Description

This message occurs when the deletion of objects and the cleanup of link failed when the operations encountered an error during a SnapMirror restore operation. Because the link could not be cleaned up

properly, the bucket could be in an inconsistent state and might not be able to accept any traffic. Also, if the links are not cleaned up, a new SnapMirror relationship cannot be established with these two buckets.

### Corrective Action

1. Run the "snapmirror delete" command to delete the link.
2. Delete the objects in the restore bucket.

### Syslog Message

Cleanup of the link between buckets "%s" and "%s" failed with error "%s" and the objects in the restore bucket could not be deleted.

### Parameters

**sourceBucket** (STRING): Name of the source bucket in the SnapMirror relationship.

**destBucket** (STRING): Name of the destination bucket in the SnapMirror relationship.

**error** (STRING): Error that was encountered.

## mgmt.snapmir.fl.rstr.lnk.cln

### Severity

ERROR

### Description

This message occurs when the cleanup of a link failed when it encountered an error during a SnapMirror restore operation. Because the link could not be cleaned up properly, the bucket could be in an inconsistent state and might not be able to accept any traffic. Also, if the links are not cleaned up, a new SnapMirror relationship cannot be established with these two buckets.

### Corrective Action

Run the "snapmirror delete" command to delete the link.

### Syslog Message

Cleanup of the link between buckets "%s" and "%s" failed with error "%s".

### Parameters

**sourceBucket** (STRING): Name of the source bucket in the SnapMirror relationship.

**destBucket** (STRING): Name of the destination bucket in the SnapMirror relationship.

**error** (STRING): Error that was encountered.

## mgmt.snapmir.init.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror@ 'initialize' or 'initialize-ls-set' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

### Syslog Message

Initialize from source volume '%s' to destination volume(s) '%s' failed with error '%s'. Job ID %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.  
**dstpaths** (STRING): Path names of the destination volumes to which the initialize failed.  
**error** (STRING): Error message.  
**jobID** (LONGINT): Job ID.

## mgmt.snapmir.ls.catchup.fail

### Severity

ALERT

### Description

This message occurs when a catch-up update for a lagging member in a load-sharing set fails.

### Corrective Action

Check the reason for the error and take action accordingly.

### Syslog Message

Load-share catch-up update transfer from source volume '%s' to destination volume '%s' failed with error '%s'. Job ID %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror® relationship.  
**dstpath** (STRING): Path name of the destination volume to which the catch-up update transfer failed.  
**error** (STRING): Error message.  
**jobID** (LONGINT): Job ID.

## mgmt.snapmir.resync.fail

### Severity

ALERT

### Description

This message occurs when a SnapMirror® 'resync' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

### Syslog Message

Resync from source volume '%s' to destination volume '%s' failed with error '%s'. Job ID %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.  
**dstpath** (STRING): Path name of the destination volume to which the resync failed.  
**error** (STRING): Error message.  
**jobID** (LONGINT): Job ID.

## mgmt.snapmir.schd.trans.fail

**Severity**

ALERT

**Description**

This message occurs when a scheduled SnapMirror® transfer fails and no automatic retry will be attempted until the next scheduled transfer.

**Corrective Action**

Check the reason for the error and take action accordingly.

**Syslog Message**

Scheduled transfer from source volume '%s' to destination volume(s) '%s' failed with error '%s'. Job ID %llu.

**Parameters**

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpaths** (STRING): Path names of the destination volumes to which the scheduled transfer failed.

**error** (STRING): Error message.

**jobID** (LONGINT): Job ID.

**mgmt.snapmir.schd.trans.overrun****Severity**

ERROR

**Description**

This message occurs when a scheduled SnapMirror® transfer overruns the schedule window.

**Corrective Action**

Check the reason for the schedule overrun and take action accordingly.

**Syslog Message**

Scheduled transfer from source volume '%s' to destination volume(s) '%s' is taking longer than the schedule window. Job ID %llu.

**Parameters**

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpaths** (STRING): Path names of the destination volumes of the scheduled transfer.

**jobID** (LONGINT): Job ID.

**mgmt.snapmir.unexpected.err****Severity**

ERROR

**Description**

This message occurs when a SnapMirror operation encounters an unexpected error. JobID might be zero if the job ID could not be determined at the time the error was generated.

**Corrective Action**

If the problem persists, contact NetApp technical support.

## Syslog Message

SnapMirror unexpected error '%s'. Job ID %llu.

## Parameters

**error** (STRING): Error message.

**jobID** (LONGINT): Job ID.

## mgmt.snapmir.update.fail

### Severity

ALERT

### Description

This message occurs when a SnapMirror® 'update' or 'update-ls-set' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

## Syslog Message

Update from source volume '%s' to destination volume(s) '%s' failed with error '%s'. Job ID %llu.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpaths** (STRING): Path names of the destination volumes to which the update failed.

**error** (STRING): Error message.

**jobID** (LONGINT): Job ID.

## mgmt.vdr events

### mgmt.vdr.fsecurity.config.err

### Severity

ERROR

### Description

This message occurs when a "snapmirror break" operation is executed while a file security job is running and the previous "snapmirror update" operation only partially transfers the changes made to the file system by the file security job. This results in the policy being applied only partially to the partner; that is, ACLs on the files are applied partially.

### Corrective Action

Verify that the file security job has been successfully applied on the Vserver DR partner by using the "job show" command on the source; it should display the "complete" status. Then resync the SnapMirror® relationship by using the "snapmirror resync -source-vserver <source-vserver-name> -destination-vserver <destination-vserver-name>" command. Alternatively, you can break the SnapMirror relationship, and then reapply the job on the destination by using the "vserver security file-directory apply -vserver <vserver-name> -policy-name <policy-name>" command. See the "vserver security file-directory" man pages for more information. If the problem persists, contact NetApp technical support.



## Syslog Message

Failed to apply file security policy "%s" on Vserver "%s".

## Parameters

**policy** (STRING): Fsecurity policy name.  
**vservername** (STRING): Name of the Vserver.

# mgmt.vopl.events

## mgmt.vopl.cleanup.failed

### Severity

ALERT

### Description

This message occurs when the system encounters an error trying to clean up after a failed 'volume move' operation.

### Corrective Action

Contact NetApp technical support..

## Syslog Message

'volume move' cleanup operation with job ID '%jd' failed in step (%s) with the error (%s) for volume '%s' in Vserver '%s'.

## Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.  
**step** (STRING): Phase during cleanup where the error occurred.  
**error** (STRING): Error recorded during the cleanup phase.  
**volName** (STRING): Name of the source volume.  
**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.del.staleLunMaps

### Severity

ERROR

### Description

This message occurs when a 'volume delete' operation cannot delete LUN initiator group maps for the LUNs in the volume, resulting in stale LUN initiator group maps. Any attempt to delete these initiator groups that are mapped to nonexistent LUNs (provided that they are not mapped to any other LUNs) fails. 'lun mapped show -igroup <igroup name>' does not list any maps, but 'igroup delete -igroup <igroup name>' fails with an error, suggesting that the initiator group cannot be deleted because it is mapped to LUNs.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Could not delete LUN initiator group maps for the volume with MSID '%llu' on Vserver '%s'. Error: %s.

## Parameters

**volMsid** (LONGINT): Master data set ID of the volume being deleted.

**vserver** (STRING): Vserver hosting the volume.

**error** (STRING): Internal software error.

## mgmt.vopl.fail.restart.nosvol

### Severity

ALERT

### Description

This message occurs when a 'volume move' operation fails in the restart phase because the source volume is not found. The system repeatedly attempts the phase until the restart can be completed. In the time between the first failure and when the restart can be completed, there might be disruption to client traffic for the volume being moved.

### Corrective Action

Verify that the job restart is completed by running the "volume move show -vserver <vservname> -volume <volName>" command. If the restart phase continues to fail, contact NetApp technical support.

### Syslog Message

'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' failed in the restart phase when attempting to retrieve the source volume.

## Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.fail.status.entry

### Severity

ALERT

### Description

This message occurs when a 'volume move' operation fails to make an entry in the VLDB table for the 'volume move show' command. The system repeatedly attempts this operation until it exhausts the number of retries. Until an entry is made, the 'volume move show' command might return stale entries use 'job show -id' to show the status of the current volume move operation.

### Corrective Action

Use the 'job show -id' command to show the status of the current 'volume move' operation. Retry the 'volume move' operation. If the problem persists, contact NetApp technical support.

### Syslog Message

The 'volume move' operation with ID '%jd' for volume name '%s' in Vserver '%s' failed while attempting to make an entry in the VLDB table for the 'volume move show' command. Use 'job show -id' command to show the current status of the job.

## Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## **mgmt.vopl.fail.status.modify**

### **Severity**

ALERT

### **Description**

This message occurs when a 'volume move' operation fails to modify the VLDB table for the 'volume move show' command. The system repeatedly attempts this phase. Until the modification succeeds, the 'volume move show' command might return stale information about the current 'volume move' operation.

### **Corrective Action**

If the problem persists, contact NetApp technical support.

### **Syslog Message**

The 'volume move' operation for volume '%s' with job ID '%jd' in Vserver '%s' cannot modify an entry in the VLDB table.

### **Parameters**

**volName** (STRING): Name of the source volume.

**jobId** (LONGINT): Job ID of the 'volume move' operation.

**vserverName** (STRING): Name of the source Vserver.

## **mgmt.vopl.made.status.entry**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a 'volume move' operation makes an entry for the 'volume move show' command because the entry did not exist in the VLDB table.

### **Corrective Action**

(None).

### **Syslog Message**

Use the 'volume move show -vserver %s -volume %s' command to show the status of this 'volume move' operation.

### **Parameters**

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## **mgmt.vopl.move.badDstVol**

### **Severity**

ERROR

### **Description**

This message occurs when a 'volume move' operation cannot access the specified destination volume.

**Corrective Action**

Verify that the destination volume exists and is online.

**Syslog Message**

The move operation could not access the specified destination volume '%s'.

**Parameters**

**dstVol** (STRING): Name of the temporary destination volume for move operation.

**mgmt.vopl.move.badSourceVol****Severity**

ERROR

**Description**

This message occurs when a 'volume move' operation cannot access the specified source volume.

**Corrective Action**

Verify that the source volume exists and is online.

**Syslog Message**

The move operation could not access the specified source volume '%s' in Vserver '%s' for job ID '%jd'.

**Parameters**

**srcVol** (STRING): Name of the volume to be moved.

**vserverName** (STRING): Name of the Vserver.

**jobID** (LONGINT): Job ID of the volume move operation.

**mgmt.vopl.move.bytesLeft****Severity**

INFORMATIONAL

**Description**

This message occurs when a 'volume move' operation is in data copy phase and we want to know the bytes left.

**Corrective Action**

(None).

**Syslog Message**

The 'volume move' operation for volume '%s' in Vserver '%s' and job ID '%jd' is in data copy phase. We want to know the number of bytes left to transfer.

**Parameters**

**volpath** (STRING): Name of the source volume of the move job.

**vserverName** (STRING): Name of the source Vserver.

**jobID** (LONGINT): Job ID of the 'volume move' operation.

## mgmt.vopl.move.cantCrtDstVol

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because the destination volume cannot be created.

### Corrective Action

Check for space on the destination aggregate. Execute the "volume move" command with the "-perform-validation-only" parameter. If it succeeds, contact NetApp technical support. If it does not succeed, increase space on the destination aggregate, then try the "volume move" command again.

### Syslog Message

Volume move job failed because destination volume "%s" could not be created.

### Parameters

**volpath** (STRING): Name of the destination volume of the move.

## mgmt.vopl.move.cantCrtSrcSnp

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because creation of a Snapshot(tm) copy on the source volume failed.

### Corrective Action

Use the "volume show" to check if the volume has enough space for the Snapshot copy. Use the "snapshot show" command to check if the number of Snapshot copies have reached the maximum count. If the issue persists, contact Contact NetApp technical support..

### Syslog Message

Volume move job failed because creation of a Snapshot copy on source volume "%s" failed.

### Parameters

**volpath** (STRING): Name of the source volume of the move job.

## mgmt.vopl.move.cantDstCncl

### Severity

NOTICE

### Description

This message occurs when the 'volume move' job cannot conclude a transfer on the destination volume.

### Corrective Action

Contact NetApp technical support.

### **Syslog Message**

'volume move' job could not conclude a transfer on the destination volume '%s'.

### **Parameters**

**volpath** (STRING): Name of the destination volume of the move job.

## **mgmt.vopl.move.cantReadCfg**

### **Severity**

ERROR

### **Description**

This message occurs when the 'volume move' job cannot read the configuration information.

### **Corrective Action**

Check whether the replicated database is online.

### **Syslog Message**

The job configuration information for 'volume move' operation with job ID '%jd' could not be read.

### **Parameters**

**jobID** (LONGINT): Job ID of the 'volume move' operation.

## **mgmt.vopl.move.cantSaveToRdb**

### **Severity**

NOTICE

### **Description**

This message occurs when the 'volume move' job cannot save information to the replicated database. The 'volume move' job attempts to retry the operation.

### **Corrective Action**

The "volume move" job is retried automatically. Check the status of database with "cluster ring show" advanced command. If the problem persists, contact NetApp technical support.

### **Syslog Message**

The volume move job '%jd' for volume '%s' in Vserver '%s' could not save the information to the replicated database.

### **Parameters**

**jobId** (LONGINT): ID of the 'volume move' job.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the Vserver.

## **mgmt.vopl.move.cantSrcCncl**

### **Severity**

NOTICE

**Description**

This message occurs when the 'volume move' job cannot conclude a transfer on the source volume.

**Corrective Action**

Verify that the source volume exists, is readable, and is not currently locked exclusively by another process. If the problem persists, contact NetApp technical support.

**Syslog Message**

'volume move' job could not conclude the transfer on the source volume '%s'.

**Parameters**

**volpath** (STRING): Name of the source volume of the move job.

**mgmt.vopl.move.ClnupFldDIVol****Severity**

ERROR

**Description**

This message occurs when a cleanup operation cannot delete the destination volume after a 'volume move' job fails.

**Corrective Action**

Delete the destination volume by using the 'volume delete' command.

**Syslog Message**

Could not delete destination volume '%s' following a failed 'volume move'.

**Parameters**

**volpath** (STRING): Name of the destination volume that could not be deleted.

**mgmt.vopl.move.cntRdDstDAttr****Severity**

NOTICE

**Description**

This message occurs when the 'volume move' job cannot read the attributes of the destination volume from disk. The 'volume move' job attempts to retry the operation.

**Corrective Action**

Use the 'volume move show' command to get the status of the current move operation. If the problem persists, contact NetApp technical support.

**Syslog Message**

The 'volume move' job could not fetch attributes of destination volume '%s' from the disk.

**Parameters**

**volpath** (STRING): Name of the destination volume involved in the move operation.

## mgmt.vopl.move.cntRdPostCutoverSrcDAttr

### Severity

ERROR

### Description

This message occurs when a 'volume move' operation cannot access the specified source volume after cutover to propagate final volume attributes. The attributes were set to values that were either obtained immediately before the cutover occurred, or were set to default values.

### Corrective Action

Verify that the attributes on the destination volume are correct. Specifically, check the value for the attribute "filesys-size-fixed".

### Syslog Message

The move operation could not read post-cutover attributes from the specified temporary source volume '%s' in Vserver '%s' for job ID '%jd'.

### Parameters

**dstVol** (STRING): Name of the temporary source volume for the move operation.

**vserverName** (STRING): Name of the Vserver.

**jobID** (LONGINT): Job ID of the volume move operation.

## mgmt.vopl.move.cntRdSrcDAttr

### Severity

NOTICE

### Description

This message occurs when the volume move operation cannot read the attributes of the source volume from disk. The volume move operation will attempt to retry the operation.

### Corrective Action

(None).

### Syslog Message

The volume move operation could not fetch the attributes of source volume '%s' from the disk.

### Parameters

**volpath** (STRING): Name of the source volume involved in the move operation.

## mgmt.vopl.move.cut.entryFail

### Severity

ERROR

### Description

This message occurs when the volume move operation cannot enter the cutover phase. The original source volume is functional. After a short delay, the volume move operation will reattempt cutover entry.



## Corrective Action

If the problem persists, reduce client traffic on the volume being moved and reduce the general load on the destination node.

## Syslog Message

The volume move operation of volume '%s' in Vserver '%s' to destination aggregate '%s' did not enter the cutover phase. The system provided the additional explanation: '%s'. After a short delay, the volume move operation will reattempt cutover entry. The job ID for the volume move job is '%jd'.

## Parameters

**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**aggrName** (STRING): Name of the destination aggregate.  
**reasonforFailure** (STRING): Reason for the failure.  
**jobID** (LONGINT): Job ID of the volume move operation.

## mgmt.vopl.move.cutover.deferred

### Severity

ALERT

### Description

This message occurs when the volume move job cannot be completed because the cutover phase has been deferred.

### Corrective Action

Address the reason for the cutover deferral, and then use the 'volume move trigger-cutover' command to reattempt the cutover phase for this volume move operation.

## Syslog Message

Cutover phase has been deferred for the volume move operation with ID '%jd' on volume '%s' on Vserver '%s' to destination aggregate '%s', because of '%s'. Space savings from efficiency could be lost.

## Parameters

**jobID** (LONGINT): Job ID of the volume move operation.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**aggrName** (STRING): Name of the destination aggregate.  
**reasonforDefer** (STRING): Reason for deferred cutover.

## mgmt.vopl.move.cutover.deferred.wait

### Severity

INFORMATIONAL

### Description

This message occurs when the volume move job cannot be completed because the cutover phase has been deferred due to the user specifying the cutover action as 'wait'.

### Corrective Action

Use the 'volume move trigger-cutover' command to attempt the cutover phase for this volume move operation.

## Syslog Message

Cutover phase has been deferred for the volume move operation with ID '%jd' on volume '%s' on Vserver '%s' to destination aggregate '%s', because of '%s'. Space savings from efficiency could be lost.

## Parameters

**jobID** (LONGINT): Job ID of the volume move operation.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**aggrName** (STRING): Name of the destination aggregate.  
**reasonforDefer** (STRING): Reason for the deferred cutover.

## mgmt.vopl.move.cutover.failed

### Severity

ERROR

### Description

This message occurs when the volume move job attempts to cut over and is unsuccessful. The system continues to retry the operation until it exhausts the number of cutover attempts allowed.

### Corrective Action

Reduce client traffic on the volume, that is scheduled to be moved.

## Syslog Message

The volume move operation with ID '%jd' on volume '%s' on Vserver '%s' to destination aggregate '%s' did not cut over because: '%s'. The volume move operation has tried to cut over '%d' out of '%d' times, %d forced.

## Parameters

**jobID** (LONGINT): Job ID of the volume move operation.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**aggrName** (STRING): Name of the aggregate.  
**reasonforFailure** (STRING): Reason for cutover failure.  
**actualcutoverAttempt** (INT): Number of actual cutover attempts made.  
**cutoverAttempt** (INT): Number of cutover attempts.  
**forcedCutoverAttempt** (INT): Number of forced cutover attempts.

## mgmt.vopl.move.cutoverMod

### Severity

INFORMATIONAL

### Description

This message occurs when a 'volume move' operation is in cutover and one of the modifiable parameters used by the 'move' job has been modified with the 'volume move modify' command. These modified values will be used for next cutover attempt.

### Corrective Action

(None).

## Syslog Message

The 'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' is performing a cutover. Values modified with the 'volume move modify' command will not be used until the next cutover attempt.

## Parameters

**jobId** (LONGINT): Job ID of the 'volume move' operation.  
**volName** (STRING): Name of the source volume.  
**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.move.done

### Severity

NOTICE

### Description

This message occurs when the volume move job succeeds.

### Corrective Action

(None).

## Syslog Message

The volume move job with ID '%jd' on volume '%s' on Vserver '%s' to destination aggregate '%s' succeeded.

## Parameters

**jobID** (LONGINT): Job ID of the volume move operation.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**aggrName** (STRING): Name of the destination aggregate.

## mgmt.vopl.move.dstPrepareFId

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because the destination volume cannot be prepared for a transfer.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Volume move job failed because the destination volume '%s' could not be prepared for a transfer.

## Parameters

**volpath** (STRING): Name of the destination volume of the move job.

## mgmt.vopl.move.findSnpDltFId

**Severity**

ERROR

**Description**

This message occurs when the volume move job cannot be completed because the call to calculate the delta between a Snapshot(tm) copy and the active file system failed.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Volume move job failed because the call to calculate the delta between a Snapshot copy and the active file system on the source volume '%s' failed.

**Parameters**

**volpath** (STRING): Name of the source volume of the move job.

**mgmt.vopl.move.getNxtSnapFld****Severity**

ERROR

**Description**

This message occurs when the 'volume move' job cannot look up the Snapshot(tm) copy next in age to the currently referenced Snapshot copy. The 'volume move' job attempts to retry the operation.

**Corrective Action**

Use the 'volume move show' command to get the status of the current move operation. If the problem persists, contact NetApp technical support.

**Syslog Message**

The 'volume move' job could not to look up the Snapshot copy next in age to the currently referenced Snapshot copy on the source volume '%s'.

**Parameters**

**volpath** (STRING): Name of the source volume involved in the move operation.

**mgmt.vopl.move.gtDstSSLstFld****Severity**

ERROR

**Description**

This message occurs when the volume move job cannot be completed because a Snapshot(tm) copy list cannot be obtained from the destination volume.

**Corrective Action**

Check whether the destination volume is online. If the problem persists, contact NetApp technical support.

**Syslog Message**

Volume move job failed because a Snapshot copy list could not be obtained from the destination volume

'%s'.

#### Parameters

**volpath** (STRING): Name of the destination volume of the move job.

### **mgmt.vopl.move.gtSnpTgSrcFld**

#### Severity

ERROR

#### Description

This message occurs when the volume move job cannot be completed because a Snapshot(tm) copy tag cannot be read from the source volume.

#### Corrective Action

Check whether the source volume is online. If the problem persists, contact NetApp technical support.

#### Syslog Message

Volume move job failed because a Snapshot copy tag could not be read from the source volume '%s'.

#### Parameters

**volpath** (STRING): Name of the source volume of the move job.

### **mgmt.vopl.move.gtSrcSSLstFld**

#### Severity

ERROR

#### Description

This message occurs when the volume move job cannot be completed because a Snapshot(tm) copy list cannot be obtained from the source volume.

#### Corrective Action

Check whether the source volume is online. If the problem persists, contact NetApp technical support.

#### Syslog Message

Volume move job failed because a Snapshot copy list could not be obtained from the source volume '%s'.

#### Parameters

**volpath** (STRING): Name of the source volume of the move job.

### **mgmt.vopl.move.gtSSTgsDstFld**

#### Severity

ERROR

#### Description

This message occurs when the volume move job cannot be completed because the Snapshot(tm) copy tags on the destination volume cannot be read.

### Corrective Action

Retry the volume move job. If the problem persists, this indicates a system level problem. Contact NetApp technical support.

### Syslog Message

Volume move job failed because Snapshot copy tags on the destination volume '%s' in vserver '%s' for the volume move of source volume '%s' in vserver '%s' could not be read.

### Parameters

**dstVol** (STRING): Name of the destination volume of the move job.

**srcVol** (STRING): Name of the source volume of the move job.

**dstVserverName** (STRING): Name of destination vserver.

**srcVserverName** (STRING): Name of source vserver.

## mgmt.vopl.move.has.clones

### Severity

ERROR

### Description

This message occurs when a 'volume move' operation finishes moving the volume to the destination aggregate but cannot delete the original volume instance on the source aggregate. This is because the original volume instance has volume clones.

### Corrective Action

The system usually reaps these types of volume after all clones have been deleted or split automatically. If space occupied by source needs to be freed immediately, then delete the original volume instance using the 'volume delete' command after all the volume clones have been deleted or split.

### Syslog Message

The 'volume move' operation with ID '%jd' could not delete the original volume %s in Vserver %s on %s aggregate following a successful volume move because the original volume instance has volume clones.

### Parameters

**jobId** (LONGINT): Job ID of the volume move operation.

**volName** (STRING): Name of the volume instance that could not be deleted.

**vserverName** (STRING): Name of the volume's Vserver.

**aggrName** (STRING): Name of the source aggregate.

## mgmt.vopl.move.invlDXfrState

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because a transfer is in an invalid state.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Volume move job with job ID '%jd' for volume '%s' in Vserver '%s' failed because a transfer was in an invalid state.

## Parameters

**jobID** (LONGINT): ID of the 'volume move' job.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.jobIgnoreQuit

### Severity

ERROR

### Description

This message occurs when a volume move job fails to stop when you issue a 'volume move abort' or 'job stop' command. If the command is not repeated, the volume move job might continue to run to completion.

### Corrective Action

Retry the 'volume move abort' or 'job stop' command before the volume move job is complete.

## Syslog Message

Volume move job '%jd' for volume '%s' in Vserver '%s' could not be stopped at this time. Retry the 'volume move abort' or 'job stop' command.

## Parameters

**jobID** (LONGINT): ID of the volume move job.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.jobQuit

### Severity

INFORMATIONAL

### Description

This message occurs when a 'volume move' job is not completed because a 'volume move abort' or 'job stop' command was issued.

### Corrective Action

Verify the status of the move job through the 'volume move show' command.

## Syslog Message

'volume move' job with ID '%jd' for volume '%s' in Vserver '%s' was not completed because a 'volume move abort' or 'job stop' command was issued.

## Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.last.upd.fail

### Severity

ERROR

### Description

This message occurs when a 'volume move' operation fails to update attributes on the destination volume after the cutover has been completed.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' failed to update the attributes on the destination volume after the cutover has been completed.

### Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.move.mark.source.reapable.failed

### Severity

ERROR

### Description

This message occurs when a 'volume move' operation finishes moving the volume to the destination aggregate but cannot mark the original volume instance on the source aggregate for deletion. The original volume instance cannot be deleted immediately because it has volume clones.

### Corrective Action

Delete the original volume instance using the 'volume delete' command after all the volume clones have been deleted or split.

### Syslog Message

The 'volume move' operation with ID '%jd' could not mark the Volume %s containing volume clones in Vserver %s on %s aggregate for later deletion after a successful volume move.

### Parameters

**jobId** (LONGINT): Job ID of the volume move operation.

**volName** (STRING): Name of the volume instance that could not be marked for later deletion.

**vserverName** (STRING): Name of the volume's Vserver.

**aggrName** (STRING): Name of the destination aggregate.

## mgmt.vopl.move.memAllocFld

### Severity

NOTICE



## Description

This message occurs when the 'volume move' job cannot allocate memory. The 'volume move' job attempts to retry the operation.

## Corrective Action

Use the 'volume move show' command to get the status of the current move operation. If the problem persists, contact NetApp technical support.

## Syslog Message

The 'volume move' job ID '%jd' for volume '%s' in Vserver '%s' could not allocate memory.

## Parameters

**jobId** (LONGINT): ID of the 'volume move' job.  
**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.moveModlgr

### Severity

INFORMATIONAL

### Description

This message occurs when modifiable parameters of the 'move' job have been modified by 'volume move modify' command, but the 'volume move' operation finished with older values and did not use the newly modified values.

### Corrective Action

Use 'volume move show' command to verify the parameters used during cutover phase and the final state of the volume move operation.

### Syslog Message

The 'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' has completed. Parameters modified during last cutover attempt through 'volume move modify' command have been ignored or 'volume move' operation failed on job restart before refreshing parameters from the VLDB.

### Parameters

**jobId** (LONGINT): Job ID of the 'volume move' operation.  
**volName** (STRING): Name of the source volume.  
**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.move.nosnapspace

### Severity

ERROR

### Description

This message occurs when a 'volume move' operation fails to create a Snapshot(tm) copy on source volume because the number of Snapshot copies in the volume exceeds the maximum limit or there is insufficient disk space in the volume.

## Corrective Action

Make sure that the Snapshot copy limit is not reached for the volume. The current Data ONTAP version allows a maximum of 255 Snapshot copies per volume. If required, delete old Snapshot copies of the volumes on the aggregate by running the 'volume snapshot delete' command.

## Syslog Message

The 'volume move' operation with job ID '%jd' for volume '%s' present on Vserver '%s' cannot proceed (Reason: Cannot create a Snapshot copy on the volume).

## Parameters

**jobId** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.nospace

### Severity

ERROR

### Description

This message occurs when there is not enough space on the volume or aggregate for a 'volume move' operation to proceed. It could be on the source or the destination node.

## Corrective Action

Reduce the amount of used space on the volume or aggregate, or increase the size of the volume or aggregate. To increase the size of the volume, run the 'volume modify' command. To create space on an aggregate, take one of the following actions: a. To free up more space on a volume, delete old Snapshot(tm) copies of the volumes on the aggregate by running the 'volume snapshot delete' command. b. To increase the size of a volume or an aggregate, add disks to the aggregate by running the 'storage aggregate add-disks' command.

## Syslog Message

The 'volume move' operation with ID '%jd' for volume '%s' present on Vserver '%s' cannot proceed (Reason: Not enough space in the volume or aggregate).

## Parameters

**jobId** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.parent.online

### Severity

NOTICE

### Description

This message occurs when a "volume move" operation fails to take offline the quiesced volume remaining after moving a clone parent, because wafliron is running.

## Corrective Action

No action is required, because the system removes the quiesced clone parent automatically. To free some of the resources used by the volume, you can manually take it offline by using the "volume offline"

command after the wafiron operation has completed.

### **Syslog Message**

The "volume move" operation for the volume %s on vserver %s with job-id %jd could not offline the quiesced clone parent volume left behind.

### **Parameters**

**volName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.  
**jobId** (LONGINT): Job ID of the "volume move" operation.

## **mgmt.vopl.move.queryDstFld**

### **Severity**

ERROR

### **Description**

This message occurs when the volume move job cannot be completed because a query of the destination for transfer status failed.

### **Corrective Action**

Its a system level problem. Contact Contact NetApp technical support. for further investigation.

### **Syslog Message**

Volume move job failed because a query for transfer status of the destination volume '%s' in vserver '%s' failed.

### **Parameters**

**volpath** (STRING): Name of the destination volume of the move job.  
**vserverName** (STRING): Name of destination vserver.

## **mgmt.vopl.move.querySrcFld**

### **Severity**

NOTICE

### **Description**

This message occurs when the 'volume move' job cannot be completed because a query of the source for transfer status failed. The 'volume move' job attempts to retry the operation.

### **Corrective Action**

Use the 'volume move show' command to get the status of the move operation. If the problem persists, contact NetApp technical support.

### **Syslog Message**

The 'volume move' job could not query the source '%s' for transfer status.

### **Parameters**

**volpath** (STRING): Name of the source volume involved in the move operation.

## mgmt.vopl.move.refVldbFld

### Severity

NOTICE

### Description

This message occurs when a 'volume move' operation fails to refresh modifiable parameters used by the 'move' job from the Volume Location Database (VLDB). Volume move operation will use current parameters until a refresh attempt succeeds.

### Corrective Action

(None).

### Syslog Message

The 'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' could not refresh the modifiable parameters used by the 'move' job from the VLDB. Volume move operation will use current parameters until a refresh attempt succeeds.

### Parameters

**jobId** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.move.restart.failed

### Severity

ERROR

### Description

This message occurs when a volume move job encounters conditions that prevent it from restarting; for example, required Snapshot(tm) copies have been deleted. This might result in loss of client access to the volume, stale temporary volumes and stale metadata.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Volume move job ID '%jd' could not restart correctly and might have left a stale temporary volume and stale persistent volume move data. Client access to the data stored on the source volume might be blocked. Reason: %s, error %d. Move job was started for volume '%s' in Vserver '%s'.

### Parameters

**jobID** (LONGINT): Job ID of the volume move job.

**reason** (STRING): Reason for restart failure.

**error** (INT): Error code for the restart failure.

**volName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.snpNtDeletd

**Severity**

ERROR

**Description**

This message occurs when a 'volume move' operation cannot delete a reference Snapshot(tm) copy created by the move job.

**Corrective Action**

Use the 'snapshot show' command to verify that the Snapshot copy exists and is not locked by another process. Use the 'snapshot delete' command to delete the Snapshot copy manually.

**Syslog Message**

Snapshot copy for 'volume move' job ID %jd could not be deleted: name: %s; Vserver ID: %d; volume name: %s.

**Parameters**

**jobId** (LONGINT): ID of the volume move job.

**snapName** (STRING): Name of the Snapshot copy.

**vserverId** (INT): Destination Vserver ID.

**volName** (STRING): Name of the parent volume of the Snapshot copy.

**mgmt.vopl.move.source.cleanup.failed****Severity**

ERROR

**Description**

This message occurs when a 'volume move' operation finishes moving the volume to the destination aggregate but cannot delete the original volume instance on the source aggregate.

**Corrective Action**

Delete the original volume instance using the 'volume delete' command after offlining the volume.

**Syslog Message**

The 'volume move' operation with ID '%jd' could not delete the Volume %s in Vserver %s on %s aggregate after a successful volume move.

**Parameters**

**jobId** (LONGINT): Job ID of the volume move operation.

**volName** (STRING): Name of the volume instance that could not be deleted

**vserverName** (STRING): Name of the volume's Vserver.

**aggrName** (STRING): Name of the destination aggregate.

**mgmt.vopl.move.srcVAttrGtFld****Severity**

ERROR

**Description**

This message occurs when the volume move job cannot be completed because the Volume Location Database (VLDB) attributes for the source volume cannot be read.

**Corrective Action**

To check whether the VLDB process is up and running, run the (privilege: advanced) "cluster ring show -unitname vldb" command. Contact NetApp technical support.

**Syslog Message**

Volume move job failed because the VLDB attributes for source volume "%s" could not be read.

**Parameters**

**volpath** (STRING): Name of the source volume of the move job.

**mgmt.vopl.move.srcVolInconst****Severity**

ERROR

**Description**

This message occurs when the volume move job cannot be completed because the source volume is inconsistent.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Volume move job failed because the source volume '%s' is inconsistent.

**Parameters**

**volpath** (STRING): Name of the source volume of the move job.

**mgmt.vopl.move.srcVolOffln****Severity**

ERROR

**Description**

This message occurs when the volume move job cannot be completed because the source volume is offline.

**Corrective Action**

Bring the source volume online and retry the move job.

**Syslog Message**

Volume move job failed because the source volume '%s' is offline.

**Parameters**

**volpath** (STRING): Name of the source volume of the move.

**mgmt.vopl.move.srcVolRstrctd****Severity**

ERROR

## Description

This message occurs when the volume move job cannot be completed because the source volume is restricted.

## Corrective Action

If the source volume is of type DP and a "snapmirror initialize" operation is running on it, wait for the operation to finish, then retry "volume move" operation. If the problem persists, contact NetApp technical support.

## Syslog Message

Volume move job failed because the source volume "%s" is restricted.

## Parameters

**volpath** (STRING): Name of the source volume of the move job.

## mgmt.vopl.move.srcVolTemp

### Severity

ALERT

### Description

This message occurs when a 'volume move' operation cannot start because the source volume is marked as temporary.

### Corrective Action

Verify that the source volume exists.

### Syslog Message

'volume move' operation with job ID '%jd' for source volume name '%s' and Vserver name '%s' failed because the source volume is marked as temporary.

### Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.move.stSnpTgDstFld

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because a Snapshot(tm) copy tag cannot be written to the destination volume.

### Corrective Action

Check whether the destination volume is online. If the problem persists, contact NetApp technical support.

### Syslog Message

Volume move job failed because a Snapshot copy tag could not be written to the destination volume '%s'.

## Parameters

**volpath** (STRING): Name of the destination volume of the move job.

## mgmt.vopl.move.voplIntrnlErr

### Severity

ERROR

### Description

This message occurs when the volume move job cannot be completed because of an internal error in the Volume Operations Library used by the job management code.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Volume move job ID '%jd' for volume '%s' in Vserver '%s' failed because of an internal error in Volume Operations Library used by the job management code.

## Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.move.XferNotFound

### Severity

ERROR

### Description

This message occurs when the 'volume move' job cannot find the specified transfer.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Transfer not found for 'volume move' job with '%jd' for volume '%s' in Vserver '%s'.

## Parameters

**jobId** (LONGINT): ID of the 'volume move' job.

**volName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

## mgmt.vopl.rehost.failed

### Severity

ERROR

### Description

This message occurs when a volume rehost job attempts to rehost a volume from one Vserver to another and is unsuccessful. This might happen if the target volume is being accessed by some other Data



ONTAP® operations (for example, volume move); if the target volume is in a relationship, such as a SnapMirror DP relationship; or if the aggregate hosting the volume is undergoing a wafliron process. The system makes its best effort to revert the rehost procedure and make the volume available under the source Vserver. Subsequent attempts to rehost this volume might be successful; if the volume rehost revert procedure still fails, use the command "debug vreport" to fix volume discrepancies.

### Corrective Action

Follow the instructions provided as part of the volume rehost error and retry the operation after some time. If the volume rehost revert procedure fails and there are volume discrepancies, use command "debug vreport" in diagnostic mode to fix volume discrepancies.

### Syslog Message

The volume rehost operation on volume '%s [%s]' on Vserver '%s' to destination Vserver '%s' by administrator '%s' failed with reason: '%s'. The volume rehost operation can be retried after sometime.

### Parameters

**volName** (STRING): Name of the target volume.

**volUuid** (STRING): Target volume UUID.

**vserverName** (STRING): Name of the source Vserver.

**destinationVserverName** (STRING): Name of the destination Vserver.

**adminName** (STRING): Name of the administrator attempting the volume rehost operation.

**reasonforFailure** (STRING): Reason for volume rehost failure.

## mgmt.vopl.rehost.succeeded

### Severity

INFORMATIONAL

### Description

This message occurs when a volume rehost job is successful in rehosting a volume from one Vserver to another.

### Corrective Action

(None).

### Syslog Message

The volume rehost operation on volume '%s [%s]' on Vserver '%s' to destination Vserver '%s' by administrator '%s' is successful.

### Parameters

**volName** (STRING): Name of the target volume.

**volUuid** (STRING): Target volume UUID.

**vserverName** (STRING): Name of the source Vserver.

**destinationVserverName** (STRING): Name of the destination Vserver.

**adminName** (STRING): Name of the administrator attempting the volume rehost operation.

## mgmt.vopl.rstart.fail.jstate

### Severity

NOTICE

## Description

This message occurs when a volume move operation encounters an error during the restart phase because it cannot retrieve the state information from the job manager. The system repeatedly attempts the restart process until it can be completed.

## Corrective Action

Verify that the restart phase is completed by running the "volume move show -vserver <vservname> -volume <volName%gt;" command. If the restart phase continues to fail, contact NetApp technical support.

## Syslog Message

volume move operation with job ID '%jd' for volume '%s' in Vserver '%s' encountered an error in the restart phase when attempting to retrieve the job state information.

## Parameters

**jobID** (LONGINT): Job ID of the volume move operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.restart.retry.fail

### Severity

ALERT

### Description

This message occurs when an attempt to restart an interrupted 'volume move' operation repeatedly encounters an error and cannot proceed any further. The 'volume move' job will transition to a failed state without cleanup being attempted.

### Corrective Action

Check the job status for the ID listed in the event message, and look for errors involving the source and destination nodes and aggregates, and then contact NetApp technical support.

### Syslog Message

The attempt to restart the 'volume move' operation with job ID '%jd' for volume '%s' in Vserver '%s' failed.

### Parameters

**jobID** (LONGINT): Job ID of the 'volume move' operation.

**volName** (STRING): Name of the source volume.

**vserverName** (STRING): Name of the source Vserver.

## mgmt.vopl.vmsi.fail.no.entry

### Severity

ALERT

### Description

This message occurs when a "volume move" operation cannot update the VLDB table, because the system cannot find any record of the operation within the table. This failure prevents the operation from being displayed in by the "volume move show" command.

**Corrective Action**

Use the "job show -id" command to display the current status of the "volume move" operation.

**Syslog Message**

The "volume move" operation for the volume %s with job-id %jd could not modify the entry in the VLDB table for the "volume move show" command. Use the "job show -id" command to track the status of this move operation.

**Parameters**

**volName** (STRING): Name of the volume.

**jobId** (LONGINT): Job ID of the "volume move" operation.

# mgmtgwd events

## mgmtgwd.certificate events

### mgmtgwd.certificate.expired

#### Severity

ERROR

#### Description

This message occurs when a digital certificate for a Vserver expires. Client-server communication is not secure.

#### Corrective Action

Install a new digital certificate on the system using the 'security certificate create' or 'security certificate install' command.

#### Syslog Message

A digital certificate with Fully Qualified Domain Name (FQDN) %s, Serial Number %s, Certificate Authority '%s' and type %s for Vserver %s has expired.

#### Parameters

**common\_name** (STRING): Name of the Fully Qualified Domain Name(FQDN) whose digital certificate expired.

**serial\_number** (STRING): Serial Number of the digital certificate.

**certificate\_authority** (STRING): Name of the Certificate Authority.

**type** (STRING): Type of the digital certificate.

**vserver** (STRING): Name of the Vserver.

### mgmtgwd.certificate.expiring

#### Severity

ERROR

#### Description

This message occurs when a digital certificate for a Vserver is about to expire. Client-server communication will not be secure if the certificate expires.

#### Corrective Action

Install a new digital certificate on the system using the 'security certificate create' or 'security certificate install' command.

#### Syslog Message

A digital certificate with Fully Qualified Domain Name (FQDN) %s, Serial Number %s, Certificate Authority '%s' and type %s for Vserver %s will expire in the next %d day(s).

#### Parameters

**common\_name** (STRING): Name of the Fully Qualified Domain Name(FQDN) whose digital certificate is about to expire.

**serial\_number** (STRING): Serial Number of the digital certificate.

**certificate\_authority** (STRING): Name of the Certificate Authority whose digital certificate is about to

expire.

**type** (STRING): Type of the digital certificate.

**vserver** (STRING): Name of the Vserver.

**days** (LONGINT): Number of days remaining until the certificate expires.

## mgmtgwd.configbr events

### mgmtgwd.configbr.backupCompleted

#### Deprecated

This event was replaced by configbr.backupCompleted.

#### Severity

NOTICE

#### Description

This message occurs when a scheduled configuration backup is successfully created.

#### Corrective Action

(None).

#### Syslog Message

Scheduled configuration backup %s was created successfully.

#### Parameters

**backup** (STRING): Backup file.

### mgmtgwd.configbr.backupFailed

#### Deprecated

This event was replaced by configbr.backupFailed.

#### Severity

ERROR

#### Description

This message occurs when when the configuration backup cannot be created.

#### Corrective Action

If the reason is about the root volume file system, ensure that it is available and has enough space to create a backup. If the reason is that a scheduled backup file could not be created, create the backup file manually and upload it to the remote URL.

#### Syslog Message

The %s configuration backup %s cannot be created. Error: '%s'.

#### Parameters

**type** (STRING): Type of backup, either cluster or local.

**backup** (STRING): Backup file.

**reason** (STRING): Reason for the error.

## mgmtgwd.configbr.deleteFailed

### Deprecated

This event was replaced by configbr.deleteFailed.

### Severity

ERROR

### Description

This message occurs when an old configuration backup cannot be deleted.

### Corrective Action

Ensure that the node is accessible and healthy. If necessary, delete the backup manually.

### Syslog Message

The %s configuration backup %s on node %s cannot be deleted. Error: '%s'.

### Parameters

**type** (STRING): Type of backup, either cluster or local.

**backup** (STRING): Backup file.

**node** (STRING): Node on which the backup file exists.

**reason** (STRING): Reason for the error.

## mgmtgwd.configbr.distributeFailed

### Deprecated

This event was replaced by configbr.distributeFailed.

### Severity

ERROR

### Description

This message occurs when a configuration backup cannot be distributed to another node in the cluster.

### Corrective Action

Ensure that the destination node is accessible and healthy. If necessary, upload the backup file to the remote URL to increase the availability of the backup file.

### Syslog Message

Configuration backup %s cannot be distributed to node %s. Error: '%s'.

### Parameters

**backup** (STRING): Backup file.

**destination\_node** (STRING): Destination node to which the backup file failed to be distributed.

**reason** (STRING): Reason for the error.

## mgmtgwd.configbr.initFailed

### Deprecated

This event was replaced by configbr.initFailed.

**Severity**

ERROR

**Description**

This message occurs when when the configuration backup and restore feature cannot initialize the necessary software components.

**Corrective Action**

If the reason is about the backup directory under the root volume, ensure that the root volume is available and has enough free space. If the reason is about populating replicated database entries, wait 5 minutes and then use the advanced mode of the command-line interface to access the "Backup Settings" menu to ensure that backup schedules are listed. If the reason is about creating backup jobs, wait 5 minutes and then run the 'job show' command to ensure that backup jobs are displayed for each schedule.

**Syslog Message**

Configuration backup and restore feature could not be initialized. Error: '%s'.

**Parameters**

**reason** (STRING): Reason for the initialization failure.

**mgmtgwd.configbr.intError****Deprecated**

This event was replaced by configbr.intError.

**Severity**

ERROR

**Description**

This message occurs when an attempt to create a configuration backup file encountered an internal software error.

**Corrective Action**

Create the configuration backup file manually by using the "system configuration backup create" command.

**Syslog Message**

Configuration backup and restore detected an internal software error. Error: '%s'.

**Parameters**

**error** (STRING): Internal software error.

**mgmtgwd.configbr.noDestURL****Deprecated**

This event was replaced by configbr.noDestURL.

**Severity**

ERROR

**Description**

This message occurs when the cluster backup is saved on only one node and no offsite configuration

backup destination URL is configured.

### Corrective Action

Configure an offsite configuration backup destination URL as soon as possible, using the commands "system configuration backup settings modify -destination <destination-url> -username <username-on-destination-url>" and "system configuration backup settings set-password".

### Syslog Message

Cluster backup is saved on only one node and no offsite configuration backup destination URL is configured.

### Parameters

(None).

## mgmtgwd.configbr.snapshotDeleteFailed

### Deprecated

This event was replaced by configbr.snapshotDelFailed.

### Severity

ERROR

### Description

This message occurs when a Snapshot(tm) copy on the root volume cannot be deleted. The Snapshot copy was created during a configuration backup.

### Corrective Action

Ensure that the root volume file system is available.

### Syslog Message

Snapshot copy %s on the root volume cannot be deleted on node %s. Error: '%s'.

### Parameters

**name** (STRING): Name of the Snapshot copy.

**node** (STRING): Node on which the Snapshot copy was created.

**reason** (STRING): Reason for the error.

## mgmtgwd.configbr.unpubFailed

### Deprecated

This event was replaced by configbr.unpubFailed.

### Severity

ERROR

### Description

This message occurs when the configuration backup and restore feature fails to unpublish a package by using the file replication service. The configuration backup process creates file packages and registers them with the file replication service, which stores the package in the root file system. When the configuration backup process is complete, it cleans up these files by unpublishing the package by using the file replication service. When unpublishing fails, these files are not be cleaned up and remain in the root file system. This message does not indicate a failure in the configuration backup process, which continues



despite the event. You can monitor the status of an ongoing configuration backup process by using the "job show" command. The completed configuration backup files are visible when you run the "system configuration backup show" command.

### Corrective Action

Obtain diagnostic privilege and view the file replication packages by using the "system services replication show" command. Unpublish the file replication packages by issuing the "system services replication unpublish" command, for which you also need diagnostic privilege. Check the "cluster ring show" command to ensure that the node is in quorum; management quorum is required for the unpublish command to succeed. Do not take these actions while an active configuration backup process is running.

### Syslog Message

Configuration backup and restore failed to unpublish package %s; Error: '%s'.

### Parameters

**package\_name** (STRING): Name of the package.  
**reason** (STRING): Reason for the unpublishing failure.

## mgmtgwd.configbr.uploadFailed

### Deprecated

This event was replaced by configbr.uploadFailed.

### Severity

ERROR

### Description

This message occurs when the configuration backup cannot be uploaded to the destination URL.

### Corrective Action

Ensure that the destination URL is reachable, the protocol in the URL is supported, and the user credentials are valid.

### Syslog Message

Configuration backup file %s cannot be uploaded to the destination URL %s. Error: '%s'.

### Parameters

**backup** (STRING): Backup file.  
**destination** (STRING): Destination URL.  
**reason** (STRING): Reason for the upload failure.

## mgmtgwd.fg events

### mgmtgwd.fg.32bit.nfs3ids

### Severity

ERROR

### Description

This message occurs when the user creates a FlexGroup on a Vserver that has NFSv3 enabled, but uses 32-bit identifiers, which could cause file ID collisions.

### Corrective Action

To prevent file ID collisions within the FlexGroup, enable NFSv3 64-bit FSIDs and file IDs on the Vserver. For more information or assistance, contact NetApp technical support.

### Syslog Message

FlexGroup %s was created on Vserver %s which uses 32-bit NFSv3 FSIDs and File IDs. File ID collisions may occur.

### Parameters

**volumeName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmtgwd.fg.32bit.nfs4ids

### Severity

ERROR

### Description

This message occurs when the user creates a FlexGroup on a Vserver that has NFSv4.x enabled, but uses 32-bit identifiers, which could cause file ID collisions.

### Corrective Action

To prevent file ID collisions within the FlexGroup, enable NFSv4.x 64-bit FSIDs and file IDs on the Vserver. For more information or assistance, contact NetApp technical support.

### Syslog Message

FlexGroup %s was created on Vserver %s which uses 32-bit NFSv4.x FSIDs and File IDs. File ID collisions may occur.

### Parameters

**volumeName** (STRING): Name of the volume.  
**vserverName** (STRING): Name of the Vserver.

## mgmtgwd.fg.modify.constituent.fail

### Severity

ERROR

### Description

This message occurs when modification of a FlexGroup constituent's state fails during the operation to modify the state of the FlexGroup. The FlexGroup will be in the 'mixed' state and constituents that are still online will be able to serve data. Constituents that are offline will not be able to serve data. All management operations will be unavailable on the FlexGroup until it is brought completely online.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact NetApp technical support.

### Syslog Message

Failed to modify the state of FlexGroup constituent %s in Vserver %s. Reason: %s.

## Parameters

**volumeName** (STRING): Name of the constituent.

**vserverName** (STRING): Name of the Vserver.

**reason** (STRING): Reason for the failure.

## mgmtgwd.fg.modify.rlbk.fail

### Severity

ERROR

### Description

This message occurs after a series of failures. First, the FlexGroup resize fails. To recover, we try to rollback the FlexGroup constituents to their original sizes. If we fail to rollback any of these constituents, this message will occur.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact NetApp technical support.

### Syslog Message

Failed to rollback the size of FlexGroup constituent %s in Vserver %s Reason: %s.

## Parameters

**volumeName** (STRING): Name of the constituent.

**vserverName** (STRING): Name of the Vserver.

**reason** (STRING): Reason for the failure.

## mgmtgwd.fg.sis.op.constituent.fail

### Severity

ERROR

### Description

This message occurs when the execution of the efficiency operation on a FlexGroup constituent fails.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact Netapp technical support for further assistance.

### Syslog Message

Failed to execute the efficiency operation on FlexGroup constituent %s in Vserver %s. Operation: %s.  
Reason: %s.

## Parameters

**volumeName** (STRING): Name of the constituent.

**vserverName** (STRING): Name of the Vserver.

**operationType** (STRING): Type of the efficiency operation.

**reason** (STRING): Reason for the failure.

# mgmtgwd.filereplication events

## mgmtgwd.filereplication.failverify

### Severity

ERROR

### Description

This message occurs when the startup scan of the file replication publishing root directory fails to verify packages. This might be due to file corruption or a mixed-version issue.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

The startup scan of the file replication publishing directory failed to verify one or more packages.

### Parameters

(None).

## mgmtgwd.filereplication.forbidden

### Severity

ERROR

### Description

This message occurs when the file replication service receives a request on an unsecure network interface.

### Corrective Action

Use the "network interface show" command to verify that the source IP address of the request is from an internal, authorized source. This could be a security breach attempt. File replication never attempts to communicate over nonclustered interfaces.

### Syslog Message

The file replication service denied an incoming request that came from '%s' through the virtual interface address '%s'.

### Parameters

**client\_ip** (STRING): IP address of the client.

**server\_ip** (STRING): IP address of the server interface.

## mgmtgwd.filereplication.nopubroot

### Severity

ERROR

### Description

This message occurs when the file replication service cannot find or initialize the publishing root directory.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

The publishing root directory at '%s' could not be initialized. Error: '%s'.

**Parameters**

**dir** (STRING): Location of the publishing root directory.

**reason** (STRING): Reason the directory cannot be used.

**mgmtgwd.filereplication.noserver****Severity**

ERROR

**Description**

This message occurs when the file replication module cannot initialize the local server process.

**Corrective Action**

Ensure that the local replication network port is not in use ("network port show" and "network port broadcast-domain show") and that system resources are not limited.

**Syslog Message**

The file replication service could not initialize its server process. Error: '%s'.

**Parameters**

**description** (STRING): Reason for the failure.

**mgmtgwd.filereplication.nostart****Severity**

ERROR

**Description**

This message occurs when the file replication service cannot be started.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

The file replication service could not be started.

**Parameters**

(None).

**mgmtgwd.fips events****mgmtgwd.fips.post.failure**

**Severity**

ERROR

**Description**

This message occurs when an unrecoverable error prevents initialization of the cryptographic subsystem into FIPS mode. No SSL communication will be possible with the storage system in this state. Other ONTAP features may also be affected by the lack of cryptographic operations.

**Corrective Action**

To operate in non FIPS cryptographic mode, turn off FIPS mode for the SSL interface using command "security config modify -interface SSL -is-fips-enabled false" and restart the system. For more information or assistance, contact NetApp technical support.

**Syslog Message**

SSL interface disabled due to failure of FIPS mode cryptographic services.

**Parameters**

(None).

## mgmtgwd.jobmgr events

### mgmtgwd.jobmgr.fetcherr

**Severity**

ERROR

**Description**

This message occurs when the system cannot locate the configuration for a specific Vserver.

**Corrective Action**

Disable and reenble the auditing function of the affected Vserver using the 'vserver audit disable' and 'vserver audit enable' commands to attempt to retrieve the missing configuration information. If the audit enable command fails, contact the system administrator for assistance.

**Syslog Message**

Job '%s' [jobid %d, name "%s", UUID %s] failed on Vserver '%s'.

**Parameters**

**name** (STRING): Name of the job.  
**jobid** (INT): ID of the job.  
**description** (STRING): Description of the job.  
**jobuuid** (STRING): UUID of the job.  
**vservername** (STRING): Vserver name linked with the job.

### mgmtgwd.jobmgr.high.jobs

**Severity**

ERROR

## Description

This message occurs when the total number of queued jobs, including both the cluster-wide jobs and the node-local jobs, exceeds a configured threshold.

## Corrective Action

Allow the backlog of jobs to run to completion before adding more jobs. Halt external tools from issuing frequent commands that result in jobs. You can view the job queue thresholds with the diag privilege command "job queue config show". The command "job show" displays cluster-wide jobs. The command "job private show" displays node-local jobs. The command "job history show" displays job history. The command "job queue show" displays the job queues.

## Syslog Message

High number of jobs: there are %d jobs queued on this node. New jobs will not be accepted by this node when the total number of jobs exceeds %d.

## Parameters

**jobs** (LONGINT): Number of jobs queued.  
**max** (LONGINT): Maximum number of jobs accepted.

## mgmtgwd.jobmgr.high.jobtype

### Severity

ERROR

### Description

This message occurs when the total number of queued jobs of any type, including both the cluster-wide jobs and the node-local jobs, exceeds a configured threshold.

### Corrective Action

Allow the backlog of jobs to run to completion before adding more jobs. Halt external tools from issuing frequent commands that result in jobs. You can view the job queue thresholds with the diag privilege command "job queue config show". The command "job show" displays cluster-wide jobs. The command "job private show" displays node-local jobs. The command "job history show" displays job history. The command "job queue show" displays the job queues.

### Syslog Message

High number of "%s" jobs: there are %d of this job type queued on this node. New jobs will not be accepted by this node when the total number of jobs of this type exceeds %d.

### Parameters

**type** (STRING): Job type.  
**jobs** (LONGINT): Number of jobs queued of this type.  
**max** (LONGINT): Maximum number of jobs accepted of this type.

## mgmtgwd.jobmgr.init.err

### Severity

ERROR

### Description

This message occurs when the Job manager cannot load a private job.

## Corrective Action

Look into the private job history by using the `job_id`. If necessary, restart the activity by using the applicable command. Example Command: `job private show-completed -node node-01 -id 10`

## Syslog Message

"Unable to load private job. Job UUID = %s and ID= %ul. Reason: %s"

## Parameters

`job_uuid` (STRING): UUID of the private job.  
`job_id` (LONGINT): Human-readable ID of the private job.  
`failure_reason` (STRING): Reason for failure to load.

## mgmtgwd.jobmgr.jobcomplete.failure

### Severity

INFORMATIONAL

### Description

This message occurs when the Job Manager completes a job unsuccessfully.

### Corrective Action

(None).

### Syslog Message

Job "%s" [id %d] (%s) completed unsuccessfully: %s (%llu).

### Parameters

`name` (STRING): Name of the job.  
`jobid` (INT): ID of the job.  
`description` (STRING): Description of the job.  
`completion` (STRING): Completion string for the job.  
`code` (LONGINT): Completion code for the job.  
`jobuuid` (STRING): UUID of the job.

## mgmtgwd.jobmgr.jobrestart

### Severity

INFORMATIONAL

### Description

This message occurs when the Job Manager restarts a job.

### Corrective Action

(None).

### Syslog Message

Restarted job "%s" [id %d]: %s.

### Parameters

`name` (STRING): Name of the job.  
`jobid` (INT): ID of the job.



**description** (STRING): Description of the job.  
**jobuuid** (STRING): UUID of the job.

## **mgmtgwd.jobmgr.private.jobcomplete.failure**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the Job Manager completes a private job unsuccessfully.

### **Corrective Action**

(None).

### **Syslog Message**

Private job "%s" [id %d] (%s) completed unsuccessfully: %s (%llu).

### **Parameters**

**name** (STRING): Name of the job.  
**jobid** (INT): ID of the job.  
**description** (STRING): Description of the job.  
**completion** (STRING): Completion string for the job.  
**code** (LONGINT): Completion code for the job.  
**jobuuid** (STRING): UUID of the job.

## **mgmtgwd.jobmgr.sched.peer.update**

### **Severity**

ERROR

### **Description**

This message occurs when a job schedule is updated on the peer cluster and might need to be updated on this cluster.

### **Corrective Action**

Update the schedule on this cluster to match the peer cluster.

### **Syslog Message**

"Schedule "%s" was %s on peer cluster and must also be updated on this cluster. Command on peer cluster: %s."

### **Parameters**

**schedule** (STRING): Schedule name.  
**operation** (STRING): Operation performed on schedule: "created," "modified," or "deleted."  
**command** (STRING): Job schedule command that was executed on the peer cluster.

## **mgmtgwd.jobmgr.sched.type**

### **Severity**

ERROR

## Description

This message occurs when a job schedule is created with the same name as a schedule from the MetroCluster(tm) partner cluster, but with a different schedule type.

## Corrective Action

Create a replacement schedule with a different name, and then delete one of the same-named schedules.

## Syslog Message

Schedule "%s" was created with the same name as an existing schedule, but with a different schedule type.

## Parameters

**schedule** (STRING): Schedule name.

## mgmtgwd.jobmgr.sched.update

### Deprecated

Deprecated as of version 9.6. Job schedules are automatically replicated to the peer cluster.

### Severity

ERROR

## Description

This message occurs when a job schedule is updated on this cluster and might need to be updated on the peer cluster.

## Corrective Action

On the peer cluster, update the schedule to match this cluster.

## Syslog Message

"Schedule "%s" was %s and must also be updated on the peer cluster. Command: %s."

## Parameters

**schedule** (STRING): Schedule name.

**operation** (STRING): Operation performed on schedule: "created," "modified," or "deleted."

**command** (STRING): Job schedule command that was executed.

## mgmtgwd.mixed events

### mgmtgwd.mixed.version

### Severity

ALERT

## Description

This message occurs when not all nodes are running the same version of ONTAP. Not all features of a software version are available to the cluster unless all nodes are at the same version. This message can occur when upgrade or revert is in progress or has not completed successfully.

## Corrective Action

Use "version" to check the current cluster version and use "version -node \*" to check the version on individual nodes.

## Syslog Message

Not all nodes are at the same ONTAP version. Cluster is operating in version %s. One or more nodes are at %s.

## Parameters

**ClusterVersion** (STRING): Current cluster version.

**HighestVersion** (STRING): Highest version in the cluster.

# mgmtgwd.replace events

## mgmtgwd.replace.node.mgmt.lif

### Severity

ERROR

### Description

This message occurs when the temporary node management LIF, which is normally deleted automatically, has not been replaced. It should be replaced to free the DHCP address.

### Corrective Action

Create a new node management LIF by using the 'network interface create' command, and then delete the temporary node management LIF by using the 'network interface delete' command.

## Syslog Message

Temporary node management LIF '%s' should be replaced with a new node management LIF.

## Parameters

**lif** (STRING): Temporary node management LIF.

# mgmtgwd.repos events

## mgmtgwd.repos.conflictingOps

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when an Infinite Volume operation cannot be started because another Infinite Volume operation is already in progress.

### Corrective Action

Check whether another Infinite Volume operation is currently in progress by running the 'job show' command. After the current operation is complete, try the new operation again. If the error persists, contact NetApp technical support.

## Syslog Message

Infinite Volume operation '%s' could not be started on volume '%s' on Vserver '%s', because operation '%s' is already in progress.

## Parameters

**attemptedOperation** (STRING): Name of the Infinite Volume operation that could not be started.

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**existingOperation** (STRING): Name of the Infinite Volume operation already running.

# mgmtgwd.revertto events

## mgmtgwd.revertTo.forceErrors

### Deprecated

This event was replaced by revertTo.force.

### Severity

NOTICE

### Description

This message occurs when a revert to process is being initiated with the "force" flag, and configuration errors have been detected.

### Corrective Action

(None).

### Syslog Message

Forcing revert in spite of detected configuration errors.

### Parameters

(None).

# mgmtgwd.rootvol events

## mgmtgwd.rootvol.recovery.changed

### Severity

EMERGENCY

### Description

This message occurs when the system detects that a root volume might have changed, causing the local management databases to be out of sync with the replicated databases. The node is not fully operational. Data lufs might have failed over within the cluster, so NFS and CIFS access is limited on this node. There also might have been a loss of stored data. Administrative capability is limited to local recovery procedures for this node.

### Corrective Action

Contact technical support to obtain the root volume recovery procedures.

## **Syslog Message**

The contents of the root volume might have changed and the local management databases might be out of sync with the replicated databases. This node is not fully operational. Contact technical support to obtain the root volume recovery procedures.

## **Parameters**

(None).

## **mgmtgwd.rootvol.recovery.different**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the system detects a different root volume than it has used previously. The node is not fully operational. Data lifs might have failed over within the cluster, so NFS and CIFS access is limited on this node. Administrative capability is limited to local recovery procedures for this node.

### **Corrective Action**

Contact technical support to obtain the root volume recovery procedures.

## **Syslog Message**

The root volume is different than the one previously used by this node. This node is not fully operational. Contact technical support to obtain the root volume recovery procedures.

## **Parameters**

(None).

## **mgmtgwd.rootvol.recovery.new**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the system detects a new root volume. The node is not fully operational. Data lifs might have failed over within the cluster, so NFS and CIFS access is limited on this node. There also might be a loss of stored data. Administrative capability is limited to local recovery procedures for this node.

### **Corrective Action**

Contact technical support to obtain the root volume recovery procedures.

## **Syslog Message**

A new root volume was detected. This node is not fully operational. Contact technical support to obtain the root volume recovery procedures.

## **Parameters**

(None).

## **mgmtgwd.rootvol.recovery.takeover.changed**

## Severity

EMERGENCY

## Description

This message occurs when the system detects that a root volume might have changed, causing the local management databases to be out of sync with the replicated databases. This occurred due to corruption of the Non-Volatile RAM log (NVLOG) during takeover. The node is not fully operational. Data logical interfaces (LIFs) might have failed over within the cluster, so NFS and CIFS access is limited on this node. There also might have been a loss of stored data. Administrative capability is limited to local recovery procedures for this node.

## Corrective Action

Contact technical support to obtain the root volume recovery procedures.

## Syslog Message

The contents of the root volume might have changed and the local management databases might be out of sync with the replicated databases due to corruption of NVLOG data during takeover. This node is not fully operational. Contact technical support to obtain the root volume recovery procedures.

## Parameters

(None).

# mgmtgwd.rootvolrec events

## mgmtgwd.rootvolrec.low.space

## Severity

EMERGENCY

## Description

This message occurs when the system detects that the root volume is dangerously low on space. The node is not fully operational. Data logical interfaces (LIFs) might have failed over within the cluster, so NFS and CIFS access is limited on this node. Administrative capability is limited to local recovery procedures for this node to make space available on the root volume.

## Corrective Action

Make space available on the root volume by deleting old Snapshot® copies, deleting unneeded files from /mroot, or expanding the root volume capacity. After space is available, reboot the controller. For more information or assistance, contact NetApp technical support.

## Syslog Message

The root volume on node "%s" is dangerously low on space. Less than %d MB of free space remaining.

## Parameters

**node** (STRING): Node on which the root volume is low on space.

**threshold\_in\_mb** (INT): Threshold, in MB, that is being used to compare against the root volume's available space.

# mgmtgwd.schedsnap events

## mgmtgwd.schedSnap.createFail

### Severity

ERROR

### Description

This message occurs when the system fails to create a Snapshot(tm) copy for the given volume.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact NetApp technical support.

### Syslog Message

Skipping creation of scheduled Snapshot copy '%s' on volume '%s' on vserver '%s' for schedule '%s' prefix '%s' (reason: %s)

### Parameters

**snapshotName** (STRING): Name of the Snapshot copy.

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the virtual server.

**scheduleName** (STRING): Name of the Snapshot policy schedule.

**prefixName** (STRING): Name of the prefix associated with the Snapshot policy schedule.

**reason** (STRING): Reason for the failure.

## mgmtgwd.schedSnap.createSkip

### Severity

ERROR

### Description

This message occurs when a scheduled Snapshot(tm) copy for a volume is skipped because there is a conflicting operation being performed on the volume.

### Corrective Action

Correct the problem based on the reason specified. If the error persists, contact NetApp technical support.

### Syslog Message

Skipping creation of scheduled Snapshot copy for volume "%s" on Vserver "%s" using policy "%s". Reason: %s.

### Parameters

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**snapshotPolicy** (STRING): Name of the Snapshot policy.

**reason** (STRING): Reason that the scheduled Snapshot copy was skipped.

**appUuid** (STRING): UUID of the Application.

# mgmtgwd.sign events

## mgmtgwd.sign.cert.expired

### Severity

ERROR

### Description

This message occurs when a digital certificate for a Vserver expires. Client-server communication is not secure.

### Corrective Action

Sign a new digital certificate on the system using the 'security certificate sign' command.

### Syslog Message

A digital certificate with Fully Qualified Domain Name (FQDN) %s, Serial Number %s, Certificate Authority '%s' and Certificate Authority Serial Number %s for Vserver %s has expired.

### Parameters

**common\_name** (STRING): Name of the Fully Qualified Domain Name(FQDN) whose digital certificate expired.

**serial\_number** (STRING): Serial Number of the digital certificate.

**certificate\_authority** (STRING): Name of the Certificate Authority.

**ca\_serial\_number** (STRING): Serial Number of the Certificate Authority(CA) that signed the digital certificate.

**vserver** (STRING): Name of the Vserver.

## mgmtgwd.sign.cert.expiring

### Severity

ERROR

### Description

This message occurs when a digital certificate for a Vserver is about to expire. Client-server communication will not be secure if the certificate expires.

### Corrective Action

Sign a new digital certificate on the system using the 'security certificate sign' command.

### Syslog Message

A digital certificate with Fully Qualified Domain Name (FQDN) %s, Serial Number %s, Certificate Authority '%s' and Certificate Authority Serial Number %s for Vserver %s will expire in %d day(s).

### Parameters

**common\_name** (STRING): Name of the Fully Qualified Domain Name(FQDN) whose digital certificate is about to expire.

**serial\_number** (STRING): Serial Number of the digital certificate.

**certificate\_authority** (STRING): Name of the Certificate Authority whose digital certificate that is about to expire.

**ca\_serial\_number** (STRING): Serial Number of Certificate Authority(CA) that signed the digital certificate.

**vserver** (STRING): Name of the Vserver.

**days** (INT): Number of days remaining until the certificate expires.



# mgmtgwd.snapshot events

## mgmtgwd.snapshot.deleteWarn

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs when a Snapshot(tm) copy cannot be deleted from an Infinite Volume.

### Corrective Action

Delete one or more older Snapshot copies associated with this schedule by using the "volume snapshot delete" command.

### Syslog Message

Skipping deletion of scheduled Snapshot copy "%s" on volume "%s" on Vserver "%s" for schedule "%s" prefix "%s" (reason: %s).

### Parameters

**snapshotName** (STRING): Name of the Snapshot copy.

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**scheduleName** (STRING): Name of the Snapshot policy schedule.

**prefixName** (STRING): Name of the prefix associated with the Snapshot policy schedule.

**reason** (STRING): Reason for the failure.

## mgmtgwd.snapshot.frcRestore

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when a Snapshot(tm) copy is restored on an Infinite Volume using the 'force' option.

### Corrective Action

(None).

### Syslog Message

Volume '%s' on Vserver '%s' was restored to Snapshot copy '%s' using the 'force' option.

### Parameters

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**snapshotName** (STRING): Name of the Snapshot copy.

## mgmtgwd.snapshot.partCreate

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when a Snapshot(tm) copy is partially created on an Infinite Volume. Residual Snapshot copies might remain on one or more Infinite Volume constituent volumes. These Snapshot copies might consume the Snapshot reserve on the constituent volumes.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Scheduled Snapshot copy '%s' of volume '%s' on Vserver '%s' for schedule '%s' prefix '%s' was partially created and cleanup failed on one or more constituent volumes.

### Parameters

**snapshotName** (STRING): Name of the Snapshot copy.

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

**scheduleName** (STRING): Name of the Snapshot policy schedule.

**prefixName** (STRING): Name of the prefix associated with the Snapshot policy schedule.

**constituentMsidList** (STRING): List of master data set IDs (MSIDs) of Infinite Volume constituent volumes that might have residual Snapshot copies.

## mgmtgwd.snapshot.tooManySnap

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when a Snapshot(tm) copy cannot be created on an Infinite Volume because too many Snapshot copies exist for the associated schedule.

### Corrective Action

Delete an older Snapshot copy associated with this schedule using the 'volume snapshot delete' command.

### Syslog Message

Scheduled Snapshot copy '%s' of volume '%s' on Vserver '%s' could not be created because too many Snapshot copies exist for the associated schedule.

### Parameters

**snapshotName** (STRING): Name of the Snapshot copy.

**volumeName** (STRING): Name of the volume.

**vserverName** (STRING): Name of the Vserver.

# mgmtgwd.upgrade events

## mgmtgwd.upgrade.aborted

### Severity

ALERT

### Description

This message occurs when an upgrade task is aborted. The node will not be fully operational.

### Corrective Action

(Call support).

### Syslog Message

(None).

### Parameters

(None).

## mgmtgwd.upgrade.failed

### Severity

ALERT

### Description

This message occurs when one or more upgrade tasks fails. The node will not be fully operational.

### Corrective Action

(Call support).

### Syslog Message

(None).

### Parameters

(None).

## mgmtgwd.upgrade.started

### Severity

INFORMATIONAL

### Description

This message occurs when an upgrade is started.

### Corrective Action

(None).

### Syslog Message

(None).

**Parameters**

(None).

**mgmtgwd.upgrade.success****Severity**

INFORMATIONAL

**Description**

This message occurs when an upgrade is completed successfully.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mgmtgwd.var events****mgmtgwd.var.recovery****Severity**

EMERGENCY

**Description**

This message occurs when the system detects a missing or corrupted /var root directory. The node is not fully operational. Data LIFs might have failed over within the cluster, so NFS and CIFS access is limited on this node. There also might be a loss of stored data. Administrative capability is limited to local recovery procedures for this node.

**Corrective Action**

Contact NetApp technical support to obtain the recovery procedures.

**Syslog Message**

Part of the /var boot directory is missing. This node is not fully operational. Contact technical support to obtain the recovery procedures.

**Parameters**

(None).

**mgmtgwd.volcreatedrestricted events****mgmtgwd.volcreatedrestricted****Severity**

NOTICE

## Description

This message occurs when a volume is created in the restricted state. This message is used to update the status of the health dashboard.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**vserver** (STRING): The vserver that owns the volume.

**volume** (STRING): The name of the volume.

# mgmtgwd.volmove events

## mgmtgwd.volmove.notoptimized

### Severity

NOTICE

### Description

This message occurs when a volume move operation does not take the optimized path to avoid read requests from a cloud-based tier.

### Corrective Action

(None).

### Syslog Message

Volume move operation from source aggregate "%s" to destination aggregate "%s" of volume "%s" on Vserver "%s" might need to copy data from existing objects in the capacity tier.

### Parameters

**sourceAggregate** (STRING): Name of the source aggregate.

**destinationAggregate** (STRING): Name of the destination aggregate.

**volume** (STRING): Name of the volume.

**Vserver** (STRING): Vserver that owns the volume.

**errMsg** (STRING): Error reason why volume move operation did not avoid reads from the cloud.

## mgmtgwd.volmove.optimized

### Severity

NOTICE

### Description

This message occurs when a volume move operation avoids a read request from a cloud-based tier.

### Corrective Action

(None).

## Syslog Message

Volume move operation from source aggregate "%s" to destination aggregate "%s" of volume "%s" on Vserver "%s" avoided copying data from existing objects in the capacity tier.

## Parameters

**sourceAggregate** (STRING): Name of the source aggregate.

**destinationAggregate** (STRING): Name of the destination aggregate.

**volume** (STRING): Name of the volume.

**Vserver** (STRING): Vserver that owns the volume.

# mgmtgwd.vreport events

## mgmtgwd.vreport.nodesUnreachable

### Severity

ERROR

### Description

This message occurs when the 'debug vreport' command found that some nodes were unreachable. Such nodes may contain WAFL-VLDB discrepancies which will not get reported.

### Corrective Action

The file mgwd.log logs the list of unavailable nodes. Ensure all these nodes are reachable from the node where the 'debug vreport' is run.

## Syslog Message

Vreport encountered some unreachable nodes. The report may be incomplete.

## Parameters

(None).

# mgmtgwd.wafl events

## mgmtgwd.wafl.rdb.autorecover

### Severity

NOTICE

### Description

This message occurs when the system will automatically enter RDB recovery because the system has detected that a root volume might have changed, possibly causing the local management databases to be out of sync with the replicated databases. This occurred due to corruption of the Non-Volatile RAM log (NVLOG) during takeover. The node is not fully operational. Data logical interfaces (LIFs) might have failed over within the cluster, so NFS and CIFS access is limited on this node. There also might have been a loss of stored data. Administrative capability is limited to local recovery procedures for this node.

### Corrective Action

Use the command 'event log show' to monitor for RDB recovery completion. On successful RDB recovery completion, the event 'rdb.recovery.succeeded' will be generated. Should RDB recovery fail, the event 'rdb.recovery.failed' will be generated. Failure will also generate a 'callhome.root.vol.recovery.reqd' event.

**Syslog Message**

The system will automatically enter RDB recovery because the system has detected that a root volume might have changed, possibly causing the local management databases to be out of sync with the replicated databases.

**Parameters**

(None).

# mgr events

## mgr.avail events

### mgr.avail.stat.reset

**Severity**

NOTICE

**Description**

This message is emitted when statistics is reset

**Corrective Action**

(None).

**Syslog Message**

Availability statistics reset

**Parameters**

(None).

## mgr.boot events

### mgr.boot.bios\_backup\_flash

**Severity**

ERROR

**Description**

This message occurs when the system boots from its backup BIOS image because either it failed to boot from its primary BIOS image within the expected timeframe or the image is corrupted.

**Corrective Action**

Power-cycle the system. If the system still does not boot from the primary BIOS image, contact technical support for assistance.

**Syslog Message**

System boot from the primary BIOS flash image failed. system booted from the backup BIOS flash image.

**Parameters**

(None).

### mgr.boot.disk\_done

**Severity**

INFORMATIONAL

**Description**

This message is sent when a system completes disk boot.



**Corrective Action**

(None).

**Syslog Message**

%s boot complete. Last disk update written at %s

**Parameters**

**ver** (STRING): (long) version of the OS which just booted

**last\_write\_time** (STRING): time at which the previous OS last wrote to disk

**mgr.boot.floppy\_25\_7****Severity**

NOTICE

**Description**

This message is sent when a system completes a 25/7 boot.

**Corrective Action**

(None).

**Syslog Message**

%s floppy boot complete (after 25/7).

**Parameters**

**ver** (STRING): (long) version of the OS which just booted

**mgr.boot.floppy\_done****Severity**

INFORMATIONAL

**Description**

This message is sent when a system completes a plain old floppy boot.

**Corrective Action**

(None).

**Syslog Message**

%s floppy boot complete.

**Parameters**

**ver** (STRING): (long) version of the OS which just booted

**mgr.boot.floppy\_media****Severity**

EMERGENCY

**Description**

This message is sent when a system completes a bypass\_media boot.

**Corrective Action**

(None).

**Syslog Message**

%s floppy boot complete (with disk media errors bypassed).

**Parameters**

**ver** (STRING): (long) version of the OS which just booted

**mgr.boot.floppy\_prev\_cp****Severity**

INFORMATIONAL

**Description**

This message is sent when a system completes a prev\_cp boot.

**Corrective Action**

(None).

**Syslog Message**

%s floppy boot complete (after prev\_cp).

**Parameters**

**ver** (STRING): (long) version of the OS which just booted

**mgr.boot.mroot.done****Severity**

INFORMATIONAL

**Description**

This event marks the end of populating the mroot during the boot sequence; this event is generated immediately after the process of populating the mroot has completed.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mgr.boot.mroot.start**

**Severity**

INFORMATIONAL

**Description**

This event marks the start of populating the mroot during the boot sequence; this event is generated immediately before the process of populating the mroot begins.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mgr.boot.new\_OS2****Severity**

NOTICE

**Description**

This message is sent when a system boots with a different OS version from before, and the previous version is known.

**Corrective Action**

(None).

**Syslog Message**

This system is running a new version of %s (was %s, now %s)

**Parameters**

**osName** (STRING): The flavor of operating system (Data ONTAP or NetCache)

**oldOsVersion** (STRING): The previous OS version

**newOsVersion** (STRING): The new OS version

**mgr.boot.node.unstable****Severity**

ALERT

**Description**

This message occurs when a node does not reach a stable state within 5 minutes after system boot. The system will continue booting while waiting for a stable state. Under a takeover scenario, the node has to be stable for giveback to be completed.

**Corrective Action**

Run the "storage failover show" command for takeover status. If the issue persists, contact NetApp technical support for assistance.

## Syslog Message

The node has not reached a stable state within 5 minutes after system boot. Retries are in progress.

## Parameters

(None).

## mgr.boot.oldOFW

### Severity

ERROR

### Description

This message occurs when the system firmware is determined to be too old.

### Corrective Action

For more information about upgrading your system firmware, check your support provider's web site or knowledgebase.

## Syslog Message

The current system firmware version %s is too old. Upgrade to the minimum version %s or later.

## Parameters

**currentFW** (STRING): Current system firmware version.

**newFW** (STRING): Desired system firmware version.

## mgr.boot.pmroot.done

### Severity

INFORMATIONAL

### Description

This event marks the end of populating the pmroot during the boot sequence; this event is generated immediately after the process of populating the pmroot has completed.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

(None).

## mgr.boot.pmroot.start

### Severity

INFORMATIONAL

### Description

This event marks the start of populating the pmroot during the boot sequence; this event is generated immediately before the process of populating the pmroot begins.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**mgr.boot.process.hung****Deprecated**

Deprecated as of ONTAP 9.1. Last used in 9.0x.

**Severity**

ALERT

**Description**

This message occurs when a system detects an un-responsive process during system boot. The system will continue booting while waiting for the process to respond.

**Corrective Action**

Contact NetApp technical support to diagnose the issue further.

**Syslog Message**

The %s process is not responding. Error: %s

**Parameters**

**process** (STRING): The process that is hung.

**error** (STRING): An error message about the problem.

**mgr.boot.reason\_abnormal****Severity**

EMERGENCY

**Description**

This message is sent when a system boots for abnormal reasons (panic, watchdog, etc). The message contains the reboot reason.

**Corrective Action**

(None).

**Syslog Message**

%s

**Parameters**

**reboot\_reason** (STRING): reboot reason

## **mgr.boot.reason\_ok**

### **Severity**

NOTICE

### **Description**

This message is sent when a system boots for normal reasons (halt/reboot command, etc). The message contains the reboot reason.

### **Corrective Action**

(None).

### **Syslog Message**

%s

### **Parameters**

**reboot\_reason** (STRING): reboot reason

## **mgr.boot.unequalDist**

### **Severity**

ERROR

### **Description**

This message occurs during system initialization when the local system finds that the number of disks that it can use to create a partitioned root aggregate is different than on the partner node.

### **Corrective Action**

To create a system with an equal number of automatically partitioned disks on the local and partner systems: 1. Make sure that any failed disks in the internal shelf are replaced. 2. Remove ownership from all disks. The system automatically assigns one half of the disks to the local system and the other half to the partner node. To remove disk ownership, boot into maintenance mode and manually remove ownership by using the "disk remove\_ownership" command.

### **Syslog Message**

Warning: Unequal number of disks will be used for auto-partitioning of the root aggregate on the local system and HA partner. The local system will use %d disks but the HA partner will use %d disks. To correct this situation, boot both controllers into maintenance mode and remove the ownership of all disks.

### **Parameters**

**n\_local** (INT): Number of partitionable disks that are used by the local system.

**n\_partner** (INT): Number of partitionable disks that are used by the partner system.

## **mgr.module events**

### **mgr.module.address**

### **Severity**

NOTICE

## Description

This occurs to inform you of where kernel modules were loaded in memory. This information helps NetApp technical support interpret the data in any mgr.stack.frameaddr events.

## Corrective Action

(None).

## Syslog Message

Module %s loaded from 0x%llx - 0x%llx.

## Parameters

**module\_name** (STRING): Name of the module.  
**module\_start** (PTR): Starting address of the module.  
**module\_end** (PTR): Ending address of the module.

# mgr.partner events

## mgr.partner.stack.saved

### Severity

NOTICE

### Description

This event is generated after a successful cluster takeover caused by a partner panic, and indicates partner stack trace information was collected from the time of the panic. Partner stack trace information is saved in the node status area of the cluster mailbox disk

### Corrective Action

Generally when a panic occurs, if autosupport is configured a NetApp case is automatically created, case E-mail notifications are sent, and a point of contact at your site is contacted by NetApp technical support. If autosupport is not configured or if you are not sure a case was automatically created, contact NetApp technical support.

### Syslog Message

Cluster takeover has saved partner panic stack trace information for logging.

### Parameters

(None).

# mgr.stack events

## mgr.stack.at

### Severity

NOTICE

### Description

This message contains the panic string.

**Corrective Action**

(None).

**Syslog Message**

Panic occurred at: %s

**Parameters**

**time** (STRING): The time string

**mgr.stack.checkpoint****Severity**

NOTICE

**Description**

Logs the last checkpoint dumpcore crossed before failing.

**Corrective Action**

(None).

**Syslog Message**

Dumpcore failure: At checkpoint %d

**Parameters**

**checkpoint** (INT): (None).

**mgr.stack.frame****Severity**

NOTICE

**Description**

The stack frame: frame number, functionname, base, offset.

**Corrective Action**

(None).

**Syslog Message**

Stack frame %2d: %s(0x%llx) + 0x%x

**Parameters**

**frame** (INT): The frame number

**sym** (STRING): The symbolic name

**base** (PTR): The base address

**offset** (INTHEX): The offset

**mgr.stack.frameaddr****Severity**

NOTICE



**Description**

The stack frame: frame number, (no function name), base, offset.

**Corrective Action**

(None).

**Syslog Message**

Stack frame %2d: 0x%llx + 0x%x

**Parameters**

**frame** (INT): The frame number

**base** (PTR): The base address

**offset** (INTHEX): The offset

**mgr.stack.framebase****Severity**

NOTICE

**Description**

The stack frame: frame number, base address.

**Corrective Action**

(None).

**Syslog Message**

Stack frame %2d: 0x%llx

**Parameters**

**frame** (INT): The frame number

**base** (PTR): The base address

**mgr.stack.longrun.proc****Severity**

NOTICE

**Description**

The name of the process that hogged the cpu too long.

**Corrective Action**

(None).

**Syslog Message**

Long running process: %s

**Parameters**

**name** (STRING): The name of the process

## **mgr.stack.noFrames**

### **Severity**

NOTICE

### **Description**

This message is emitted when no stack frames are available. This is normal on platforms which do not support stack tracing, or when the prior release did not support stack trace.

### **Corrective Action**

(None).

### **Syslog Message**

No stack frames available

### **Parameters**

(None).

## **mgr.stack.proc**

### **Severity**

NOTICE

### **Description**

The name of the process that we took the panic in.

### **Corrective Action**

(None).

### **Syslog Message**

Panic in process: %s

### **Parameters**

**name** (STRING): The name of the process

## **mgr.stack.saved**

### **Severity**

NOTICE

### **Description**

This message indicates that we've saved panic information in the NV RAM.

### **Corrective Action**

(None).

### **Syslog Message**

Reboot with saved panic information in log file

### **Parameters**

(None).

## **mgr.stack.string**

### **Severity**

NOTICE

### **Description**

This message contains the panic string.

### **Corrective Action**

(None).

### **Syslog Message**

Panic string: %s

### **Parameters**

**string** (STRING): The panic string

# mgwd events

## mgwd.cluscheck events

### mgwd.cluscheck.rmtDevRepair

#### Severity

NOTICE

#### Description

This message occurs when the system changes the cluster LIFs used for internal replicated database (RDB) operations, ensuring they are connected to different cluster switches for optimal cluster resiliency.

#### Corrective Action

(None).

#### Syslog Message

Cluster LIF "%s" with IP address %s connected to cluster switch "%s" is being replaced with cluster LIF "%s" with IP address %s connected to cluster switch "%s" in the RDB.

#### Parameters

**existinglif** (STRING): Cluster LIF being replaced.

**existingaddress** (STRING): IP address of the cluster LIF that is being replaced.

**existingremotedevice** (STRING): Cluster switch connected to the home port of the IP address being replaced.

**replacementlif** (STRING): Cluster LIF being added.

**replacementaddress** (STRING): IP address of the cluster LIF that is being added.

**replacementremotedevice** (STRING): Cluster switch connected to the home port of the IP address being added.

## mgwd.notify events

### mgwd.notify.halt.result

#### Severity

INFORMATIONAL

#### Description

This message reports the result of the Management Gateway process notification attempt to the Connectivity, Liveliness, Availability Monitor (CLAM) instance on its high-availability (HA) partner node that this node is undergoing a planned shutdown.

#### Corrective Action

(None).

#### Syslog Message

MGWD %s to notify CLAM on its HA partner node that this node is undergoing a planned shutdown (reason: %s). Error: %s

## Parameters

**result** (STRING): Result of the notify attempt.

**reason** (STRING): Reason for the partner beginning the halt process.

**error** (STRING): Error encountered in the notify attempt (if any).

# mhost events

## mhost.ca events

### mhost.ca.connect.cert.error

**Severity**

ERROR

**Description**

This message occurs when the client connection has failed due to certificate error.

**Corrective Action**

Ensure that a valid certificate is loaded onto the remote service to which ONTAP® is attempting to connect and the corresponding CA certificate is loaded in ONTAP Use "security certificate show" CLI command to ensure a valid CA certificate is installed for the destination service. Use "security certificate install" CLI command to install a valid CA certificate.

**Syslog Message**

Cluster agent connection: %s has failed due to certificate error. Error: %s.

**Parameters**

**client\_name** (STRING): Name of the client connection whose connection has an error.

**error\_msg** (STRING): Error that has been reported for client connection.

### mhost.ca.connect.delete

**Severity**

NOTICE

**Description**

This message occurs when the cluster agent connection is deleted upon user request.

**Corrective Action**

(None).

**Syslog Message**

Cluster agent connection to the client: %s is deleted upon user request.

**Parameters**

**client\_name** (STRING): Name of the client connection whose connection is deleted.

### mhost.ca.connect.failure

**Severity**

ERROR

**Description**

This message occurs when the cluster agent connection is not in a connected state. Reconnect attempt to establish a connection will be initiated automatically.

### **Corrective Action**

Verify that the connection path to the destination is in healthy state and network security polices are configured as needed to permit connectivity. If this behavior persists, contact NetApp technical support.

### **Syslog Message**

Cluster agent connection of the client: %s is not healthy. Attempting to reconnect. Error: %s.

### **Parameters**

**client\_name** (STRING): Name of the client connection whose connection has failed and reconnection is being attempted.

**error\_msg** (STRING): Error that has been reported for client connection.

## **mhost.ca.connect.success**

### **Severity**

NOTICE

### **Description**

This message occurs when the cluster agent connection is successfully established.

### **Corrective Action**

(None).

### **Syslog Message**

Cluster agent connection succeeded to the client: %s.

### **Parameters**

**client\_name** (STRING): Name of the client connection to which connection is established.

## **mhost.ca.msg.delivery.fail**

### **Severity**

ERROR

### **Description**

This message occurs when the client is not able to deliver the message even when the client connection is in a connected state.

### **Corrective Action**

For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "mhost.ca.msg.delivery.fail" keyword. If issue is not resolved, contact Contact NetApp technical support. for assistance.

### **Syslog Message**

Cluster agent connection: %s is not able to send messages to remote service. Error: %s.

### **Parameters**

**client\_name** (STRING): Name of the client connection whose message delivery failed.

**error\_case** (STRING): Use case which caused this error.

## mhost.ca.node.offline

### Severity

NOTICE

### Description

This message occurs when the node that is hosting the cluster agent connection is offline. ONTAP® will automatically attempt to migrate the connection to an available healthy node.

### Corrective Action

(None).

### Syslog Message

Node "%s" hosting cluster agent connection "%s", is offline. ONTAP is attempting to migrate the connection to an available healthy node.

### Parameters

**node** (STRING): Name of the node which is offline.

**client\_name** (STRING): Name of the client connection whose connection will be rehosted.

## mhost.ca.rehost.success

### Severity

NOTICE

### Description

This message occurs when the cluster agent connection of the client is rehosted successfully on a healthy node.

### Corrective Action

(None).

### Syslog Message

Cluster agent connection "%s" has migrated to node "%s".

### Parameters

**client\_name** (STRING): Name of the client connection that will be rehosted.

**node** (STRING): Name of the node to which the connection has migrated.



# migrate events

## migrate.root events

### migrate.root.completed

#### Severity

NOTICE

#### Description

This message occurs when the migration of a root aggregate has completed successfully.

#### Corrective Action

(None).

#### Syslog Message

Root aggregate migration completed successfully on node %s.

#### Parameters

**node** (STRING): Name of the node that completed the migration.

### migrate.root.failed

#### Severity

ERROR

#### Description

This message occurs when the migration of a root aggregate has failed.

#### Corrective Action

Run the "system node run -node nodename -command aggr status" command to determine if the target aggregate named "new\_root" has been created. If not, address the error condition and attempt the "system node migrate-root" command again. If the new aggregate exists, contact technical support to complete the operation.

#### Syslog Message

Root aggregate migration failed on node %s. Reason: %s.

#### Parameters

**node** (STRING): Name of the node that failed to complete the migration.

**reason** (STRING): Reason that the migrate operation failed.

### migrate.root.started

#### Severity

NOTICE

#### Description

This message occurs when the user initiates an attempt to move the root aggregate to a different set of disks on a node.

**Corrective Action**

(None).

**Syslog Message**

Root aggregate migration started on node %s.

**Parameters**

**node** (STRING): Name of the node.

# misconfig events

## misconfig.lif events

### misconfig.lif.info.lost

#### Severity

NOTICE

#### Description

This message occurs when the record for the incorrectly placed LIF is lost during MetroCluster™ switchover or switchback. As a result, the "metrocluster check lif show" command will not be able to fetch the incorrectly placed LIF on the surviving cluster after executing "metrocluster switchover" and on the disaster cluster after executing "metrocluster switchback". Also, the incorrectly placed LIF cannot be placed correctly with the "metrocluster check lif repair-placement" command.

#### Corrective Action

The LIF is incorrectly placed and needs to be manually corrected. For an incorrectly placed LIF on HA partners, correct the "home-node" by using the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -home-node <new\_home\_node>" command. For a LIF with a wrong port assignment, correct the "home-port" by using the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -home-port <new\_home\_port>" command. For a SAN LIF, use the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -status-admin down" command to bring down status-admin of the LIF, before modifying the "home-node" or the "home-port" of the LIF.

#### Syslog Message

LIF "%s" in Vserver "%s" incorrectly placed needs manual recovery. Reason: "%s".

#### Parameters

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

**error** (STRING): Reason for misconfigured LIF.

# mlm events

## mlm.arr events

### mlm.arr.affinity.not.queued

**Severity**

NOTICE

**Description**

This message occurs when the current processing of an error event prevents the Asymmetric Logical Unit Access (ALUA) handler from queuing an affinity switch event.

**Corrective Action**

(None).

**Syslog Message**

Affinity event %s not queued for LUN %s.

**Parameters**

**event** (STRING): Affinity event for asymmetric array.  
**serialno** (STRING): Serial number of the array LUN.

### mlm.arr.many.affinity.error

**Severity**

ERROR

**Description**

This message occurs when the Asymmetric Logical Unit Access (ALUA) handler encounters too many affinity errors, which might lead to reduced I/O throughput.

**Corrective Action**

Evaluate the back-end array connectivity for optimal LUN balance across array controllers. Also check for failures reported by the back-end array with regards to ALUA target port group settings. Verify that FlexArray connectivity is done as per best-practice guide.

**Syslog Message**

Failing cmd %s for LUN %s due to too many affinity errors.

**Parameters**

**cdb** (STRING): String containing the opcode byte of the command descriptor block (CDB) that was sent to the device.  
**serialno** (STRING): Serial number of the array LUN.

## mlm.array events

## mlm.array.aluaSupported

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array. Defines whether the array supports ALUA(asymmetric logical unit access) or not.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

Unsupported array: [Vendor '%s' Product '%s'] ALUA Support for array %s.

### Parameters

**vendor** (STRING): Vendor ID of the array LUN.

**product** (STRING): Product ID of the array LUN.

**alua** (STRING): Whether the array supports ALUA or not.

## mlm.array.apf.cancel

### Severity

NOTICE

### Description

This message occurs when any one of the path to an Array comes back online when the last path failure to an Array is extended.

### Corrective Action

(None).

### Syslog Message

Delay on last path to Array '%s' is cancelled as target port: WWNN: %s, WWPN: %s seen from initiator '%s' is up.

### Parameters

**arrayName** (STRING): Name of the Array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

**adapterName** (STRING): Name of the Fibre Channel adapter accessing the array target port.

## mlm.array.apf.enter

### Severity

ALERT

### Description

This message occurs when the all-path-fail-delay option for an Array is set to a non-zero value and when last path to an Array goes down. This allows Data ONTAP to suspend I/O operations for a longer period of time before declaring a data access disruption, allowing for I/O operations to resume if any path comes

back online within the specified duration.

### Corrective Action

In such cases, you should evaluate the specified device for possible replacement. Sometimes the problem results from the partial failure of a component involved in the Fibre Channel. Common things to check include properly secured cable connections, and crimped or otherwise damaged cables. If the adapter is in use check cabling.

### Syslog Message

Last path to Array '%s' target port: WWNN: %s, WWPN: %s seen from initiator '%s' is down. Failing this path will be extended for %d seconds.

### Parameters

**arrayName** (STRING): Name of the Array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

**adapterName** (STRING): Name of the Fibre Channel adapter accessing the array target port.

**delay** (INT): All Path Fail delay in seconds.

## mlm.array.apf.exit

### Severity

NOTICE

### Description

This message occurs when all path to an Array goes down and the Array remains inaccessible even after the extended duration is set by all-path-fail-delay option.

### Corrective Action

(None).

### Syslog Message

No path to Array '%s' came back online when last path failure is extended to %d seconds.

### Parameters

**arrayName** (STRING): Name of the Array.

**delay** (INT): All Path Fail delay in seconds.

## mlm.array.discovery

### Severity

INFORMATIONAL

### Description

This message occurs when the system discovers a new array.

### Corrective Action

(None).

### Syslog Message

A new array (%s) was connected to this node.

## Parameters

**arrayName** (STRING): Name of the array.

## mlm.array.inquiryPage80

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

```
Array LUN %s [S/N '%s' vendor '%s' product '%s'] Inquiry page 80: devtype '%d' page_code '%d' reserved '%d' page_len '%d' data '%s' hex_dump '%s'
```

## Parameters

**diskName** (STRING): Name of the array LUN.  
**serialno** (STRING): Serial number of the array LUN.  
**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**dev\_type** (INT): Device type of the array LUN.  
**page\_code** (INT): Page code of the array LUN.  
**reserved** (INT): Reserved fields in inquiry page 80.  
**page\_len** (INT): Inquiry page 80 length.  
**data** (STRING): Serial number from inquiry page 80.  
**P\_80\_hex** (STRING): Inquiry page 80 hexadecimal dump.

## mlm.array.inquiryPage83

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

```
[Vendor '%s' Product '%s'] Inquiry page 83: devtype '%d' page_code '%d' reserved '%d' page_len '%d' code_set '%d' association_id '%d' reserved_msb '%d' id_len '%d' hex_dump '%s'.
```

## Parameters

**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**dev\_type** (INT): Device type of the array LUN.  
**page\_code** (INT): Page code of the array LUN.

**reserved** (INT): Reserved fields in inquiry page 83.  
**page\_len** (INT): Inquiry page 83 length.  
**code\_set** (INT): Inquiry page 83 code set.  
**association\_id** (INT): Inquiry page 83 association ID.  
**reserved\_msb** (INT): Most significant byte of page\_len in SPC3 spec.  
**id\_len** (INT): Inquiry page 83 ID length.  
**P\_83\_hex** (STRING): Inquiry page 83 hexadecimal dump.

## mlm.array.inquiryStd

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

```
[S/N '%s' vendor '%s' product '%s'] STD inquiry : devtype '%d' dev_modifier '%d' support_level '%d'  
dev_mode '%d' revision '%s' hex_dump '%s'
```

### Parameters

**serialno** (STRING): Serial number of the array LUN.  
**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**dev\_type** (INT): Device type of the array LUN.  
**dev\_modifier** (INT): Device modifier of the array LUN.  
**support\_level** (INT): Support level.  
**dev\_mode** (INT): Development mode.  
**revision** (STRING): Product revision level(product firmware release number).  
**inquiry\_hex** (STRING): Inquiry hexadecimal dump.

## mlm.array.lun.discovery

### Severity

INFORMATIONAL

### Description

This message occurs when the system discovers a new array LUN.

### Corrective Action

(None).

### Syslog Message

```
New array LUN %s (S/N: %s) was discovered on initiator port %s at array target port WWNN: %s, WWPN:  
%s.
```



## Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**adapter** (STRING): Name of the initiator port discovering the array LUN.

**wwnn** (STRING): World Wide Node Name of the array target port exporting the new array LUN.

**wwpn** (STRING): World Wide Port Name of the array target port exporting the new array LUN.

## mlm.array.lun.forceResume

### Severity

ERROR

### Description

This message occurs when the system is forced to resume I/O to a previously quiesced array LUN path when no other viable non-quiesced paths remain. For arrays with asymmetric logical unit access a forced resume will also happen when there are no other viable non-quiesced paths on the current target port group.

### Corrective Action

Verify the configuration and check why array lun is usable through only one path.

### Syslog Message

Array LUN '%s' (%s) has had I/O resumed on the path from initiator %s to %s's target port: WWNN: %s, WWPn: %s to ensure LUN access through atleast one path. User settings are preempted.

## Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**adapterName** (STRING): Name of the initiator port used to access the array LUN.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

## mlm.array.lun.hiddenDMXVCMDB

### Severity

ERROR

### Description

This message occurs when the system determines that a Symmetrix LUN is write protected.

### Corrective Action

Either write enable the LUN or, if this is the Symmetrix VCMDB control device, consider unmapping the LUN to Data ONTAP.

### Syslog Message

Symmetrix LUN (%s, %s) is write protected and will be ignored. This logical unit may be the VCMDB control device.

## Parameters

**diskName** (STRING): Name of the LUN.

**serialno** (STRING): Serial number of the LUN.

## **mlm.array.lun.multipleIDs**

### **Severity**

ALERT

### **Description**

This message occurs when a logical device (LDEV) on an array has more than one LUN ID.

### **Corrective Action**

Verify the LUN configuration and make sure that the LUN ID is the same on all array target ports.

### **Syslog Message**

Array LUN '%s' (%s): Multiple LUN IDs are associated with the same LDEV on array %s. An LDEV can only have one LUN ID associated with it across all array target ports.

### **Parameters**

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**arrayName** (STRING): Name of the array.

## **mlm.array.lun.OnePath**

### **Severity**

ERROR

### **Description**

This message occurs when an array LUN has only one path.

### **Corrective Action**

Verify the configuration and make sure that there are at least two paths to this array LUN.

### **Syslog Message**

Array LUN '%s' is connected using only one path. To avoid the potential for data loss caused by the remaining path's failure, all array LUNs should have at least two paths.

### **Parameters**

**diskName** (STRING): Name of the array LUN.

## **mlm.array.lun.PathAdded**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a new path is added for an array LUN.

### **Corrective Action**

(None).

### **Syslog Message**

Array LUN '%s' (%s) is now connected on initiator port %s through %s's target port: WWNN: %s, WWPN:

%s at LUN ID: %d.

## Parameters

**diskName** (STRING): Name of the array LUN  
**serialno** (STRING): Serial number of the array LUN.  
**adapterName** (STRING): Name of the Fibre Channel adapter used to access the array LUN.  
**arrayName** (STRING): Name of the array.  
**wwnn** (STRING): World Wide Node Name of the array target port.  
**wwpn** (STRING): World Wide Port Name of the array target port.  
**LUNID** (INT): LUN ID to which a path was added.

## mlm.array.lun.PathFO

### Severity

NOTICE

### Description

This message occurs when an array LUN path becomes inaccessible and after a user-configurable delay, a path failover begins if the path remains inaccessible. The delay time can be modified using the "storage array modify -name array\_name -path-failover-time delay\_time" command.

### Corrective Action

(None).

### Syslog Message

Path via adapter %s and device "%s" to array "%s" is currently offline, but path failover is delayed for %d seconds. After this delay, if the path is back online, I/O will resume on this same path. Otherwise, I/O will fail over to an alternate path.

### Parameters

**adapterName** (STRING): Name of the adapter involved in the path offline operation.  
**deviceName** (STRING): Name of the Fibre Channel device.  
**arrayName** (STRING): Name of the array.  
**delay** (INT): Path failover delay, in seconds.

## mlm.array.lun.PathRemoved

### Severity

INFORMATIONAL

### Description

This message occurs when a path to this array LUN is removed.

### Corrective Action

(None).

### Syslog Message

Array LUN '%s' (%s) is no longer available on initiator port %s through %s's target port: WWNN: %s, WWPN: %s.

## Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the array LUN.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

## mlm.array.lun.quiesced

### Severity

INFORMATIONAL

### Description

This message occurs when I/O is quiesced on an array LUN.

### Corrective Action

(None).

### Syslog Message

Array LUN '%s' (%s) has had I/O quiesced on the path from initiator %s to %s's target port: WWNN: %s, WWPN: %s.

## Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the array LUN.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

## mlm.array.lun.removed

### Severity

NOTICE

### Description

This message occurs when the last path to an array LUN is removed.

### Corrective Action

(None).

### Syslog Message

Array LUN '%s' (%s) is no longer being presented to this node.

## Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

## mlm.array.lun.resumed

### Severity

INFORMATIONAL

### Description

This message occurs when I/O is resumed on an array LUN that was quiesced.

### Corrective Action

(None).

### Syslog Message

Array LUN '%s' (%s) has had I/O resumed on the path from initiator %s to %s's target port: WWNN: %s, WWPN: %s.

### Parameters

**diskName** (STRING): Name of the array LUN.

**serialno** (STRING): Serial number of the array LUN.

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the array LUN.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

## mlm.array.lun.type.control

### Severity

ERROR

### Description

This message occurs when the array LUN is a control LUN on the array. A control LUN cannot be used as a data LUN.

### Corrective Action

Disable or mask off control LUNs.

### Syslog Message

Array LUN '%s' is a control LUN and is not usable for data storage. Disable it or mask it off.

### Parameters

**diskName** (STRING): Name of the array LUN.

## mlm.array.mixedFailoverMode

### Severity

ALERT

### Description

This message occurs when Data ONTAP® determines that the set of paths to a storage device is configured with more than one failover mode policy.

## Corrective Action

Reconfigure the storage array so that each target port reports the same failover mode policy. When the reconfiguration is complete, reboot the system node that detected the error.

## Syslog Message

Fibre Channel host bus adapter %s detects that storage array %s reports a failover mode policy (%s) on target port WWNN: %s, WWPN: %s that doesn't match the policy (%s) currently in effect. The failover mode policy must be consistent on each path presented.

## Parameters

**adapterName** (STRING): Name of the Fibre Channel adapter detecting the failover mode conflict.

**arrayName** (STRING): Name of the array.

**policy1** (STRING): Inconsistent failover mode policy detected.

**wwnn** (STRING): World Wide Node Name of the array target port reporting the inconsistent failover mode policy.

**wwpn** (STRING): World Wide Port Name of the array target port reporting the inconsistent failover mode policy.

**policy2** (STRING): Failover mode policy currently in effect.

## mlm.array.ndu.complete

### Severity

INFORMATIONAL

### Description

This message occurs when you set the is-upgrade-pending flag status for the array to "false".

### Corrective Action

(None).

### Syslog Message

Array '%s' has the is-upgrade-pending flag status set to "false".

### Parameters

**arrayName** (STRING): Name of the array.

## mlm.array.ndu.doneTgtPort

### Severity

INFORMATIONAL

### Description

This message occurs when you set the is-upgrade-pending flag status for the array to "false". This message is logged for every array target port seen by the initiators from this node.

### Corrective Action

(None).

### Syslog Message

Array '%s' has the is-upgrade-pending flag status set to "false" for the path from initiator '%s' to the target port: WWNN: %s, WWPN: %s.

## Parameters

**arrayName** (STRING): Name of the array.

**adapterName** (STRING): Name of the Fibre Channel adapter accessing the array target port.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

## mlm.array.ndu.pending

### Severity

INFORMATIONAL

### Description

This message occurs when you set the is-upgrade-pending flag status to "true" for a specific array, to indicate that the array will undergo a firmware upgrade in the near future. After the upgrade is complete, you must manually set the is-upgrade-pending flag status back to "false" to maintain optimal path failover detection.

### Corrective Action

(None).

### Syslog Message

Array '%s' has its is-upgrade-pending flag status set to "true". Path failover performance will be degraded during the array firmware upgrade. After the upgrade is complete, set the is-upgrade-pending flag status for this array to "false".

## Parameters

**arrayName** (STRING): Name of the array.

## mlm.array.ndu.pendingTgtPort

### Severity

INFORMATIONAL

### Description

This message occurs when you set the is-upgrade-pending flag status for the array to "true". This message is logged for every array target port seen by the initiators from this node.

### Corrective Action

(None).

### Syslog Message

Array '%s' has the is-upgrade-pending flag status set to "true" for the path from initiator %s to the target port: WWNN: %s, WWPN: %s.

## Parameters

**arrayName** (STRING): Name of the array.

**adapterName** (STRING): Name of the Fibre Channel adapter accessing the array target port.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

## **mlm.array.ndu.portsDown**

### **Severity**

NOTICE

### **Description**

This message occurs when the array target port does not come back online after 80 seconds. This message is logged for every array target port seen by the initiators of this node that does not come back online after the array firmware upgrade.

### **Corrective Action**

(None).

### **Syslog Message**

Array '%s' target port: WWNN: %s, WWPN: %s seen from initiator '%s' is down.

### **Parameters**

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

**adapterName** (STRING): Name of the Fibre Channel adapter accessing the array target port.

## **mlm.array.ndu.portsOffline**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when all array ports for a specific array go offline while the array has its is-upgrade-pending flag status set to "true".

### **Corrective Action**

(None).

### **Syslog Message**

Array '%s' has target ports changed to offline state while its is-upgrade-pending status flag is set to "true".

### **Parameters**

**arrayName** (STRING): Name of the array.

## **mlm.array.ndu.portsOnline**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the first array target port is detected back online after a firmware upgrade where the array is marked with the is-upgrade-pending flag status set to "true".

### **Corrective Action**

(None).



## Syslog Message

Array '%s' has target ports changed to online state during a firmware upgrade.

## Parameters

**arrayName** (STRING): Name of the array.

## mlm.array.ndu.stillPending

### Severity

ERROR

### Description

This message occurs once every 60 minutes while the is-upgrade-pending flag status is set to "true" on an array. You should set the is-upgrade-pending flag status to "false" after the array firmware upgrade is complete to avoid path failover performance issues.

### Corrective Action

After the array firmware upgrade is complete, set the is-upgrade-pending flag status for the array to "false".

## Syslog Message

Array '%s' has the is-upgrade-pending flag status set to "true". After the array's firmware update is complete, set the is-upgrade-pending flag status for this array to "false".

## Parameters

**arrayName** (STRING): Name of the array.

## mlm.array.path.autoQuiesced

### Severity

NOTICE

### Description

This message occurs when the path to the array is quiesced, auto path quiesce happens when the number of path failures reached path\_failure count in the given duration(in mins).

### Corrective Action

(None).

## Syslog Message

Path from initiator %s to array %s target\_port:WWNN:%s,WWPN:%s has quiesced as path-failure has reached %d in %d mins duration.

## Parameters

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the path.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

**path\_failure** (INT): Number of failures occurred in the specified path.

**duration** (INT): Duration in minutes to monitor the path failures.

## **mIm.array.path.discovery**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the system discovers a new path to an array.

### **Corrective Action**

(None).

### **Syslog Message**

A new path was discovered: initiator port %s is now connected to array %s's target port: WWNN: %s, WWPN: %s. There are %d array LUNs visible on this target port.

### **Parameters**

**adapter** (STRING): Initiator port on which a new path to the array was discovered.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's newly discovered target port.

**wwpn** (STRING): World Wide Port Name of the array's newly discovered target port.

**count** (INT): Number of array LUNs visible on this target port.

## **mIm.array.path.forceResume**

### **Severity**

NOTICE

### **Description**

This message occurs when the path to the array is resumed forcefully when no other path except the quiesced path available for I/O in the node.

### **Corrective Action**

(None).

### **Syslog Message**

Path from initiator %s to array %s target\_port:WWNN:%s,WWPN:%s has resumed forcefully as it is the last available path for the I/O. User settings are preempted.

### **Parameters**

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the path.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

## **mIm.array.path.quiesced**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the path to the array is quiesced when user request the particular path to be quiesced.

**Corrective Action**

(None).

**Syslog Message**

Path from initiator %s to array %s target\_port:WWNN:%s,WWPN:%s has quiesced.

**Parameters**

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the path.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

**mlm.array.path.resumed****Severity**

INFORMATIONAL

**Description**

This message occurs when the path to the array is resumed.

**Corrective Action**

(None).

**Syslog Message**

Path from initiator %s to array %s target\_port:WWNN:%s,WWPN:%s has resumed.

**Parameters**

**adapterName** (STRING): Name of the Fibre Channel adapter used to access the path.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port.

**wwpn** (STRING): World Wide Port Name of the array's target port.

**mlm.array.port.discovery****Severity**

INFORMATIONAL

**Description**

This message occurs when the system discovers a new target port for an array.

**Corrective Action**

(None).

**Syslog Message**

A new array target port was discovered on array %s: WWNN: %s, WWPN: %s.

**Parameters**

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the new array target port.

**wwpn** (STRING): World Wide Port Name of the new array target port.

## **mlm.array.portDuplicateID**

### **Severity**

ALERT

### **Description**

This message occurs when more than one logical device (LDEV) has the same LUN ID on the same array. The error is set on these logical devices and these devices will not function until the error is resolved.

### **Corrective Action**

Ensure that each LDEV has a different LUN ID.

### **Syslog Message**

%s: %d appears more than once on %s's target port: WWNN: %s, WWPN: %s. LUN ID reuse on an array target port is not supported, even in different host groups.

### **Parameters**

**diskName** (STRING): Name of the array LUN.

**lunid** (INT): LUN ID that represents more than one LDEV.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port that is presenting this LUN ID.

**wwpn** (STRING): World Wide Port Name of the array's target port that is presenting this LUN ID.

## **mlm.array.portEmpty.direct**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a host group on an array is empty.

### **Corrective Action**

(None).

### **Syslog Message**

Fibre Channel host bus adapter %s reports no LUNs presented by storage array %s, target port WWNN: %s WWPN: %s.

### **Parameters**

**adapterName** (STRING): Name of the initiator port detecting the empty host group.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

## **mlm.array.portEmpty.switch**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a host group on an array is empty.

## Corrective Action

(None).

## Syslog Message

Fibre Channel host bus adapter %s reports no LUNs presented by storage array %, target port WWNN: %s WWPN: %s, switch:port %s.

## Parameters

**adapterName** (STRING): Name of the initiator port detecting the empty host group.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

**switchPort** (STRING): Name of the switch and port connected to the storage array.

## mlm.array.portMixedAddress

### Severity

EMERGENCY

### Description

This message occurs when a back-end array is set to Volume Set Addressing mode on one port but not on another. This results in LUN IDs that do not match in value or format. The error is set on these logical devices and these devices will not function until the error is resolved.

### Corrective Action

Consult the back-end array documentation to disable Volume Set Addressing mode on all ports.

### Syslog Message

Array %s is using inconsistent LUN addressing schemes. Check LUN addressing settings on port WWNN: %s, WWPN %s.

### Parameters

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array's target port that is reporting Volume Set Addressing Enabled.

**wwpn** (STRING): World Wide Port Name of the array's target port that is reporting Volume Set Addressing Enabled.

## mlm.array.portOneController

### Severity

ERROR

### Description

This message occurs when all paths for an array LUN are connected to the same controller on the array. This configuration results in a single point of failure at the array.

### Corrective Action

Reconfigure your back-end SAN so that array LUNs are presented from controllers in different fault domains on the target array.

## Syslog Message

All paths for array LUN %s are connected to the same controller on array %s. This configuration is a single point of failure. Ensure that array LUNs are accessible through controllers in different fault domains.

## Parameters

**diskName** (STRING): Name of the array LUN.  
**arrayName** (STRING): Name of the array.

## mlm.array.prCapabilities

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

## Syslog Message

[Vendor '%s' Product '%s'] PR\_CAP : sip\_c : %x, atp\_c : %x, res\_c : %x, ptpl\_c : %x.

## Parameters

**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**sip\_c** (INTHEX): Specify initiator ports capable bit.  
**atp\_c** (INTHEX): All target ports capable bit.  
**res\_c** (INTHEX): Reserved bit.  
**ptpl\_c** (INTHEX): Persist through power loss capable bit.

## mlm.array.prReadReservation

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

## Syslog Message

[Vendor '%s' Product '%s'] Data\_key : '0x%llx' Reservation : generation '%u' length '%u' resv\_key '0x%llx'  
scope\_address '%u' scope\_type '0x%x' extent\_len '%d' hex\_dump '%s'.

## Parameters

**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**resv\_key** (LONGINTHEX): SCSI persistent reservation key.

**generation** (INT): SCSI persistent reserve data generation value.  
**len** (INT): SCSI persistent reserve data header generation length.  
**read\_in\_resv\_key** (LONGINTEX): SCSI persistent reserve data reservation key.  
**scope\_address** (INT): SCSI persistent reserve data scope address.  
**scope\_type** (INTEX): SCSI persistent reserve data scope.  
**extent\_length** (INT): SCSI persistent reserve data extent length.  
**pr\_reservation\_hex** (STRING): SCSI persistent reservation hexadecimal dump.

## mlm.array.reportluns

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

Unsupported array: [Vendor '%s' Product '%s'] Report LUNs: lun\_count '%d' reserved field '%d'.

### Parameters

**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**lun\_count** (INT): Total number of LUNs visible.  
**reserved** (INT): Report luns reserved field.

## mlm.array.rptpgs

### Severity

ERROR

### Description

This message occurs when the system detects an array LUN from an unsupported array.

### Corrective Action

Detach the unsupported array. See the NetApp Support Site Interoperability Matrix Tool for further information about supported arrays.

### Syslog Message

Unsupported array: [Vendor '%s' Product '%s'] ASS : %d , Access flag : %d.

### Parameters

**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.  
**aas** (INT): Asymmetric Access State.  
**flags** (INT): Access state flags.

## **mlm.array.setPortTPQD**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when you set the target port queue depth for a specific target port of an array.

### **Corrective Action**

(None).

### **Syslog Message**

Array '%s' target port: WWNN: %s, WWPN: %s has a target port queue depth value set to %d.

### **Parameters**

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name of the array target port.

**wwpn** (STRING): World Wide Port Name of the array target port.

**tpqd** (INT): Target port queue depth value set for the specified target port of an array.

## **mlm.array.setTPQD**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when you set the target port queue depth for an array.

### **Corrective Action**

(None).

### **Syslog Message**

Target ports on the array '%s' have a target port queue depth value set to %d.

### **Parameters**

**arrayName** (STRING): Name of the array.

**tpqd** (INT): Target port queue depth value set for the target ports of the array.

## **mlm.array.unknown**

### **Severity**

ERROR

### **Description**

This message occurs when an array LUN is detected from an unknown array that is not supported in this version of Data ONTAP®.

### **Corrective Action**

Remove the array from the system.



## Syslog Message

Array LUN %s [S/N '%s' vendor '%s' product '%s'] is unknown and is not supported in this version of Data ONTAP.

## Parameters

**diskName** (STRING): Name of the array LUN.  
**serialno** (STRING): Serial number of the array LUN.  
**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.

## mlm.array.unsupported

### Severity

NOTICE

### Description

This message occurs when an array LUN is detected from an unsupported array.

### Corrective Action

(None).

## Syslog Message

Array LUN %s [S/N '%s' vendor '%s' product '%s'] is not supported in this version of Data ONTAP.

## Parameters

**diskName** (STRING): Name of the array LUN.  
**serialno** (STRING): Serial number of the array LUN.  
**vendor** (STRING): Vendor ID of the array LUN.  
**product** (STRING): Product ID of the array LUN.

## mlm.arraylun events

### mlm.arrayLUN.limit.exceeded

### Severity

ALERT

### Description

This message occurs when Data ONTAP® cannot install one or more discovered array LUNs since the disk device count is at full capacity.

### Corrective Action

If possible, remove an existing disk or array LUN so that array LUN installation can proceed. If the failure persists, contact NetApp technical support.

## Syslog Message

Unable to install one or more array LUNs [ vendor '%s' product '%s' ] discovered on initiator port %s, array target port WWNN: %s, WWPN: %s. The disk device limit (%d) has been reached.

## Parameters

**vendor** (STRING): Vendor ID of the storage array.  
**product** (STRING): Product ID of the storage array.  
**adapter** (STRING): Name of the initiator port discovering the array LUN(s).  
**wwnn** (STRING): World Wide Node Name of the array target port exporting the array LUN(s).  
**wwpn** (STRING): World Wide Port Name of the array target port exporting the array LUN(s).  
**limit** (INT): Disk device maximum count.

# mlm.bad events

## mlm.bad.topology.data

### Severity

ERROR

### Description

This message occurs when a device reports invalid data for "REPORT LUNS" command. The invalid data may include one or more instances of duplicate LUN IDs and/or the "list length value" from the "REPORT LUNS" command may not match the amount of data received.

### Corrective Action

Repair or replace the faulty hardware associated with the wwnn and wwpn

### Syslog Message

The device with wwnn %s and wwpn %s is reporting inconsistent "REPORT LUNS" data.

## Parameters

**wwnn** (STRING): World Wide Node Name(wwnn) of the Device.  
**wwpn** (STRING): World Wide Port Name(wwpn) of the Device.

# mlm.excessive events

## mlm.excessive.TPlatency

### Severity

NOTICE

### Description

This message occurs when ONTAP® detects that the average SCSI I/O command latency on the target port of the array exceeds the set threshold. This message indicates a degradation of I/O throughput for the target port of the array.

### Corrective Action

Reduce the I/O load on the device, increase the number of array target ports available for I/O, and/or alleviate the cause for the excessive latency on the device.

### Syslog Message

Average latency of %dus on target port %s, has exceeded %dus

## Parameters

- avg\_latency** (INT): Current average latency of commands for the array target port in microseconds.
- wwpn** (STRING): World Wide Port Name of the array target port on which the excessive latency was detected.
- threshold** (INT): Threshold value for excessive average latency in microseconds.

# mlm.non events

## mlm.non.optimized.TPusage

### Severity

ERROR

### Description

This message occurs when Data ONTAP® detects non-optimized usage of an array's target port. This message indicates degradation of I/O throughput for the indicated array's target port and the cause should be corrected.

### Corrective Action

Increase the target port queue depth for the array's target port and/or increase the number of array target ports available for I/O to alleviate the SCSI I/O command load for the array's target port. For example, to increase the maximum allowed I/O commands for an array's target port, the following command can be run: `storage array port modify -name ARRAY_NAME_1 -wwnn 50014380025d1500 -wwpn 50014380025d1508 -max-queue-depth 750`

### Syslog Message

Detected non-optimized usage of a array's target port. WWPN: %s, average service time: %dus, average latency: %dus

### Parameters

- wwpn** (STRING): World Wide Port Name of the array's target port which the non-optimized usage was detected on.
- avg\_service** (INT): Current average service time of I/O commands for the array's target port in microseconds.
- avg\_latency** (INT): Current average latency of I/O commands for the array's target port in microseconds.

# mlm.prregistration events

## mlm.PRregistration.failed

### Severity

ERROR

### Description

This message occurs when the registration of a SCSI persistent reservation key fails. A number of conditions can cause this SCSI command to fail and most are transitory in nature. If there is not a subsequent `EMS_mlm_PRregistration_succeeded` message, it is indicative of a persistent hardware fault.

### Corrective Action

A SCSI persistent registration failure normally indicates an error at the device during execution of the requested command. Such cases are often the result of an intermittent device hardware that is

automatically handled by the Data ONTAP® drivers through retries. In cases of repeated events, the specified device should be evaluated for proper operation and possible repair or replacement.

### Syslog Message

A SCSI persistent registration command failed on device %s on initiator port %s, target port World Wide Node Name %s, target port World Wide Port Name %s with a ha status of 0x%x, target status of 0x%x, sense key of 0x%x, additional sense code of 0x%x, additional sense code qualifier of 0x%x.

### Parameters

**device\_name** (STRING): Electrical path name of the device on which the error occurred.  
**ha\_name** (STRING): Name of the initiator port.  
**wwnn** (STRING): World Wide Node Name of the target port.  
**wwpn** (STRING): World Wide Port Name of the target port.  
**ha\_status** (INHEX): Integer value of the hardware adapter status for the failed command.  
**target\_status** (INHEX): Integer value of the target status for the failed command.  
**sense\_key** (INHEX): Integer value of the sense key for the failed command.  
**asc** (INHEX): Integer value of the additional sense code for the failed command.  
**asq** (INHEX): Integer value of the additional sense code qualifier for the failed command.

## mlm.PRregistration.succeeded

### Severity

INFORMATIONAL

### Description

This message occurs when a successful registration of a SCSI persistent reservation key occurs that had previously failed.

### Corrective Action

(None).

### Syslog Message

A SCSI persistent reservation command succeeded on device %s on initiator port %s, target port World Wide Node Name %s, target port World Wide Port Name %s where it had previously failed.

### Parameters

**device\_name** (STRING): Electrical path name of the device on which the registration succeeded.  
**ha\_name** (STRING): Name of the initiator port.  
**wwnn** (STRING): World Wide Node Name of the target port.  
**wwpn** (STRING): World Wide Port Name of the target port.

## mlm.scsi events

### mlm.SCSI.supportLvl.bad

### Severity

ALERT

### Description

This message occurs when an array target port contains SCSI INQUIRY data with an incorrect support level. To prevent a possible data corruption scenario, the target port is not used.

## Corrective Action

Verify that the SPC-2 bit in the array's host group is set.

## Syslog Message

The target port containing SCSI INQUIRY DATA %s:%s with a (WWPN) of %s, has an incorrect SCSI support level; it will not be available.

## Parameters

**vendorId** (STRING): Vendor ID from the SCSI INQUIRY data.

**productId** (STRING): Product ID from the SCSI INQUIRY data.

**wwpn** (STRING): World Wide Port Name (WWPN) of the target port that has the incorrect SCSI support level.

## mlm.SCSI.supportLvl.good

### Severity

NOTICE

### Description

This message occurs when an array target port that previously contained SCSI INQUIRY data with an incorrect support level has been corrected.

### Corrective Action

(None).

### Syslog Message

The target port containing SCSI INQUIRY DATA %s:%s with a (WWPN) of %s, now has a correct SCSI support level.

### Parameters

**vendorId** (STRING): Vendor ID from the SCSI INQUIRY data.

**productId** (STRING): Product ID from the SCSI INQUIRY data.

**wwpn** (STRING): World Wide Port Name (WWPN) of the target port that now has the correct SCSI support level.

## mlm.tapemc events

### mlm.tapemc.lun.discovery

### Severity

INFORMATIONAL

### Description

This message occurs when the system discovers a new tape or media changer logical unit (Tape/MC LUN) reported by a tape library.

### Corrective Action

(None).

## Syslog Message

New %s LUN %s (S/N: %s) was discovered on initiator port %s at tape library target port WWNN: %s, WWPN: %s.

## Parameters

**deviceType** (STRING): Type of LUN device.

**deviceName** (STRING): Name of the Tape/MC LUN.

**serialno** (STRING): Serial number of the LUN.

**adapterName** (STRING): Name of the initiator port discovering the LUN.

**wwnn** (STRING): World Wide Node Name(WWNN) of the library target port exporting the new LUN.

**wwpn** (STRING): World Wide Port Name(WWPN) of the library target port exporting the new LUN.

## mlm.tapemc.lun.nolnit

### Severity

ERROR

### Description

This message occurs when the initialization sequence for a tape or media changer logical unit (LUN) fails.

### Corrective Action

An expected LUN might be missing from the configuration. Information about the path to the LUN might be unknown but the product information in the log may help identify the missing LUN. Check the device LUN configuration settings.

## Syslog Message

%s device %sL%d failed LUN initialization sequence at phase %s.

## Parameters

**deviceType** (STRING): Type of LUN device.

**deviceName** (STRING): Name of the device on which the LUN initialization fails.

**lunid** (INT): LUN ID of the device on which LUN initialization fails.

**lunInitPhase** (STRING): Failing LUN initialization phase.

## mlm.tapemc.lun.PathAdded

### Severity

INFORMATIONAL

### Description

This message occurs when a new path is added for a tape or media changer logical unit (Tape/MC LUN).

### Corrective Action

(None).

## Syslog Message

%s LUN '%s' (S/N: %s) is now connected on initiator port %s through %s's target port: WWNN: %s, WWPN: %s at LUN ID: %d.

## Parameters

**deviceType** (STRING): Type of LUN device.

**deviceName** (STRING): Name of the Tape/MC LUN.  
**serialno** (STRING): Serial number of the LUN.  
**adapterName** (STRING): Name of the adapter used to access the LUN.  
**libraryName** (STRING): Name of the library.  
**wwnn** (STRING): World Wide Node Name(WWNN) of the library target port.  
**wwpn** (STRING): World Wide Port Name(WWPN) of the library target port.  
**LUNID** (INT): LUN ID to which a path was added.

## **mlm.tapemc.lun.PathRemoved**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a path to this tape or media changer logical unit (Tape/MC LUN) is removed.

### **Corrective Action**

(None).

### **Syslog Message**

The %s LUN '%s' (S/N: %s) is no longer available on initiator port %s through %s's target port: WWNN: %s, WWPN: %s.

### **Parameters**

**deviceType** (STRING): Type of LUN device.  
**deviceName** (STRING): Name of the Tape/MC LUN.  
**serialno** (STRING): Serial number of the LUN.  
**adapterName** (STRING): Name of the adapter used to access the LUN.  
**libraryName** (STRING): Name of the library.  
**wwnn** (STRING): World Wide Node Name(WWNN) of the library target port.  
**wwpn** (STRING): World Wide Port Name(WWPN) of the library target port.

## **mlm.tapemc.lun.removed**

### **Severity**

NOTICE

### **Description**

This message occurs when the last path to a tape or media changer logical unit (Tape/MC LUN) is removed.

### **Corrective Action**

(None).

### **Syslog Message**

The %s LUN '%s' (S/N: %s) is no longer visible to this node.

### **Parameters**

**deviceType** (STRING): Type of LUN device.  
**deviceName** (STRING): Name of the Tape/MC LUN.  
**serialno** (STRING): Serial number of the LUN.

## mlm.tapemc.path.discovery

### Severity

INFORMATIONAL

### Description

This message occurs when the system discovers a new tape or media changer logical unit (Tape/MC LUN) path.

### Corrective Action

(None).

### Syslog Message

A new Tape/MC path was discovered: initiator port %s is now connected to tape library %s's target port: WWNN: %s, WWPN: %s. There are %d LUNs visible on this target port.

### Parameters

**adapterName** (STRING): Name of the initiator port discovering the Tape/MC LUN.

**libraryName** (STRING): Name of the tape library.

**wwnn** (STRING): World Wide Node Name(WWNN) of the library target port exporting the new LUN.

**wwpn** (STRING): World Wide Port Name(WWPN) of the library target port exporting the new LUN.

**count** (INT): Number of LUNs visible on this port.

## mlm.tpqd events

### mlm.TPQD.TPutilization

### Severity

ERROR

### Description

This message occurs when Data ONTAP® has lowered the number of I/O commands that are allowed to be active by at least 25 percent for an array's target port. This message indicates degradation of I/O throughput for the indicated array's target port and the cause should be corrected.

### Corrective Action

Analyze the EMS logs for EMS\_fci\_device\_timeout, EMS\_scsi\_cmd\_adapterResources or EMS\_scsi\_cmd\_targetStatus indicating of "Busy" or "Queue Full" events associated to the array's target port and correct any hardware faults. Reduce the shared load on the device and/or reduce the maximum allowed I/O commands for the array's target port. For example, to reduce the maximum allowed I/O commands for an array's target port, the following command can be run: storage array port modify -name ARRAY\_NAME\_1 -wwnn 50014380025d1500 -wwpn 50014380025d1508 -max-queue-depth 256

### Syslog Message

Array target port of %s has had its dynamic target port queue depth lowered to %d which is 25%% or lower of its maximum allowed value of %d.

### Parameters

**wwpn** (STRING): World Wide Port Name of the array's target port which has its queue depth decreased by 25% or more.

**current\_qdepth** (INT): Current number of allowed active I/O commands for the array's target port.

**max\_qdepth** (INT): Maximum allowed value of the array's target port queue depth.



# mlm.vmdisk events

## mlm.vmdisk.discovery

### Severity

INFORMATIONAL

### Description

This message occurs when the system discovers a new virtual disk in a virtual machine (VM).

### Corrective Action

(None).

### Syslog Message

A new virtual disk with virtual device ID %d was discovered on initiator port %s on array %s with target port WWNN %s and WWPN: %s.

### Parameters

**VirtualDeviceID** (INT): Virtual device ID of the virtual disk.

**adapterName** (STRING): Name of the initiator port discovering the virtual disk.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name (WWNN) of the array target port exporting the new virtual disk.

**wwpn** (STRING): World Wide Port Name (WWPN) of the array target port exporting the new virtual disk.

## mlm.vmdisk.PathAdded

### Severity

INFORMATIONAL

### Description

This message occurs when a new path is added for a virtual disk in a virtual machine (VM).

### Corrective Action

(None).

### Syslog Message

Virtual disk '%s' (%s) is now connected on initiator port %s through %s's target port: WWNN: %s, WWPN: %s at virtual device ID: %d.

### Parameters

**diskName** (STRING): Name of the virtual disk

**serialno** (STRING): Serial number of the virtual disk.

**adapterName** (STRING): Name of the initiator port discovering the virtual disk path.

**arrayName** (STRING): Name of the array.

**wwnn** (STRING): World Wide Node Name (WWNN) of the array target port.

**wwpn** (STRING): World Wide Port Name (WWPN) of the array target port.

**VirtualDeviceID** (INT): Virtual device ID of the virtual disk.

# monitor events

## monitor.brokendisk events

### monitor.brokenDisk.notice

#### Severity

NOTICE

#### Description

This message occurs when a node is in a degraded state because of a broken disk. The timeout values indicate how long the system will run after each boot before shutting down because of this condition. Shutdown is to encourage you to replace broken disks. \* The timeout[123] values are duplicated as a convenience for the Syslog message, which prints out the value three times.

#### Corrective Action

Consult other logged messages to determine the identity of broken disks, or use the "storage disk show -broken" command. Replace any broken disks. The %d hour timeout can be increased by altering the "raid.timeout" value using the "storage raid-options modify" command.

#### Syslog Message

%s, the system shuts down automatically every %d hours to encourage you to replace the disk. If you reboot the system, it will run for another %d hours before shutting down.

#### Parameters

**type** (STRING): Type of degradation.

**timeout1** (INT): Timeout value.

**timeout2** (INT): Timeout value.

**timeout3** (INT): Timeout value.

## monitor.chassisfan events

### monitor.chassisFan.degraded

#### Severity

ALERT

#### Description

This message occurs when a chassis fan is degraded.

#### Corrective Action

Replace the fan unit.

#### Syslog Message

Chassis fan is degraded: %s

#### Parameters

**fan\_name** (STRING): Name of the fan.

## **monitor.chassisFan.ok**

### **Severity**

NOTICE

### **Description**

This message occurs when the chassis fans are OK.

### **Corrective Action**

(None).

### **Syslog Message**

Chassis fan %s is ok.

### **Parameters**

**fru\_name** (STRING): Name of the field-replaceable unit (FRU) containing the fan.

## **monitor.chassisFan.removed**

### **Severity**

ALERT

### **Description**

This message occurs when a chassis fan is removed This is a warning message.

### **Corrective Action**

Reinsert the FRU.

### **Syslog Message**

Chassis fan %s is removed

### **Parameters**

**fan\_name** (STRING): Name of the field-replaceable unit (FRU) containing the fan.

## **monitor.chassisFan.stop**

### **Severity**

ERROR

### **Description**

This message occurs when a chassis fan is stopped. This is a warning message.

### **Corrective Action**

Reseat the FRU. If the FRU is an I/O module, schedule unit downtime and power off the unit before reseating the FRU. If the error persists, replace the FRU.

### **Syslog Message**

Chassis fan contains at least one stopped fan: %s

### **Parameters**

**fru\_name** (STRING): Name of the field-replaceable unit (FRU) containing the fan.

## monitor.chassisFan.warning

### Severity

ALERT

### Description

This message occurs when a chassis fan is spinning either too slowly or too fast. This is a warning message.

### Corrective Action

Replace the FRU.

### Syslog Message

Chassis fan is in warning state: %s

### Parameters

**fru\_name** (STRING): Name of the field-replaceable unit (FRU) containing the fan.

## monitor.chassisfanfail events

### monitor.chassisFanFail.xMinShutdown

### Severity

EMERGENCY

### Description

This message occurs when multiple chassis fans have failed, and the system will shut down in a few minutes unless corrected.

### Corrective Action

Reseat the chassis fans. If they do not restart, replace them.

### Syslog Message

Multiple Chassis Fan failure: System will shut down in %d minutes.

### Parameters

**num\_minutes** (INT): Number of minutes until the system will shut down.

## monitor.chassispower events

### monitor.chassisPower.degraded

### Severity

ALERT

### Description

This message occurs when a power supply is degraded.

### **Corrective Action**

Degraded power might be caused by bad power supplies, bad wall power, or bad components on the motherboard. If spare power supplies are available, try exchanging them to see if that alleviates the problem. Otherwise contact NetApp technical support for further direction.

### **Syslog Message**

Chassis power is degraded: %s

### **Parameters**

**reasonText** (STRING): Description of the degradation.

## **monitor.chassisPower.ok**

### **Severity**

NOTICE

### **Description**

This message occurs when the motherboard power is OK.

### **Corrective Action**

(None).

### **Syslog Message**

Chassis power is OK.

### **Parameters**

(None).

## **monitor.chassispowersupplies events**

### **monitor.chassisPowerSupplies.ok**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when all power supplies are OK.

### **Corrective Action**

(None).

### **Syslog Message**

Chassis power supplies OK.

### **Parameters**

(None).

## **monitor.chassispowersupply events**

## **monitor.chassisPowerSupply.degraded**

### **Severity**

NOTICE

### **Description**

This message occurs when a power supply is degraded.

### **Corrective Action**

A replacement power supply might be required. Contact NetApp technical support for further direction.

### **Syslog Message**

Chassis power supply %d is degraded: %s

### **Parameters**

**ps\_number** (INT): Power supply number.

**reasonText** (STRING): Description of the degradation.

## **monitor.chassisPowerSupply.notPresent**

### **Severity**

NOTICE

### **Description**

This message occurs when an expected a power supply is not present.

### **Corrective Action**

Add a power supply to the controller chassis.

### **Syslog Message**

Chassis power supply %d is not present.

### **Parameters**

**ps\_number** (INT): Power supply number.

## **monitor.chassisPowerSupply.off**

### **Severity**

NOTICE

### **Description**

This message occurs when a power supply is off.

### **Corrective Action**

Turn on the power supply.

### **Syslog Message**

Chassis power supply %d off.

**Parameters**

**ps\_number** (INT): Power supply number.

**monitor.chassisPowerSupply.ok****Severity**

INFORMATIONAL

**Description**

This message occurs when the power supply is OK.

**Corrective Action**

(None).

**Syslog Message**

Chassis power supply %d is OK.

**Parameters**

**ps\_number** (INT): Power supply number.

**monitor.chassistemperature events****monitor.chassisTemperature.cool****Severity**

ALERT

**Description**

This message occurs when the chassis temperature is too cool. This is a warning message.

**Corrective Action**

The appliance cannot function in an environment that is too cold; find ways to warm the appliance.

**Syslog Message**

Chassis temperature is too cool: %s

**Parameters**

**describe\_toocool** (STRING): Description of the condition.

**monitor.chassisTemperature.ok****Severity**

NOTICE

**Description**

This message occurs when the chassis temperature is normal.

**Corrective Action**

(None).

## Syslog Message

Chassis temperature is ok.

## Parameters

(None).

## monitor.chassisTemperature.state.unknown

### Severity

ALERT

### Description

This message occurs when the system cannot read multiple chassis temperature sensors. The condition can be caused by software, firmware, or hardware problems.

### Corrective Action

Arrange to AC power-cycle the unit to address possible software and firmware issues. If the condition persists, contact NetApp technical support for assistance with Processor Control Module (PCM) replacement.

## Syslog Message

Chassis temperature state is unknown: %s.

## Parameters

**reasonText** (STRING): Description of the temperature sensor issue.

## monitor.chassisTemperature.warm

### Severity

ALERT

### Description

This message occurs when the chassis temperature is too warm. This is a warning message. If the condition persists, the system will shut down automatically.

### Corrective Action

Ensure that sufficient cooling air is being supplied to the chassis and that the air inlets and outlets of the unit are not blocked. If you cannot resolve the situation, shut down all clients that are attached to the storage appliance, and then shut down the system.

## Syslog Message

Chassis temperature is too warm: %s

## Parameters

**describe\_toowarm** (STRING): Description of the condition.

## monitor.completed events



## monitor.completed.iron

### Severity

ERROR

### Description

This message occurs when wafiron in the optional commit mode has finished on the aggregate but changes have not been committed or rejected. Either commit or reject the changes.

### Corrective Action

Run the "aggr wafiron commit VOL" command to commit the changes or the "aggr wafiron reject VOL" command to reject the changes.

### Syslog Message

Wafiron in the optional commit mode has completed on aggregate '%s' but the changes have not been committed or rejected. Either commit or reject the changes.

### Parameters

**volname** (STRING): Volume name

## monitor.cpu events

### monitor.cpu.ok

### Severity

INFORMATIONAL

### Description

This message occurs when the CPU busy situation is corrected.

### Corrective Action

(None).

### Syslog Message

%s

### Parameters

**report** (STRING): Description of the current CPU usage.

### monitor.cpu.tooBusy

### Severity

NOTICE

### Description

This message occurs when the CPU is too busy.

### Corrective Action

(None).

## Syslog Message

Warning: %s

## Parameters

**report** (STRING): Description of the problem.

# monitor.disklabelcheckfailed events

## monitor.diskLabelCheckFailed

### Severity

NOTICE

### Description

This message occurs when the system detects that a file system disk has an invalid or a missing RAID label. The system corrects the situation during the next label update.

### Corrective Action

No user action is required.

### Syslog Message

Periodic check of RAID %s has failed. The system will correct the problem.

### Parameters

**disk\_info** (STRING): Formatted information of the disk object.

**label\_block\_num** (INT): Label block number code.

**label\_state** (INT): Label state code.

**label\_read\_error** (INT): Label read error code.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk in RPM.

# monitor.downrevdisks events

## monitor.downRevDisks

### Severity

ERROR

### Description

This message occurs when ONTAP® detects that downrev disks are still present in the system.

### Corrective Action

Download the latest disk firmware on each node and ensure that automatic background firmware update is enabled. If necessary, use the (advanced privilege) "storage disk firmware update" command to manually

update disk firmware. Consult the man page for cautions.

### **Syslog Message**

Downrev disks still present!!

### **Parameters**

(None).

## **monitor.extcache events**

### **monitor.extCache.failed**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the monitor detects the WAFL® external cache subsystem (FlexScale) has failed and is no longer available for use.

#### **Corrective Action**

Consult the system logs to determine the original cause of the error.

#### **Syslog Message**

FlexScale WAFL external cache is unavailable due to an earlier fatal error.

#### **Parameters**

(None).

## **monitor.fan events**

### **monitor.fan.critical**

#### **Severity**

EMERGENCY

#### **Description**

This message occurs when one or more main unit fans have failed. The system cannot continue to operate.

#### **Corrective Action**

Replace the failed fans.

#### **Syslog Message**

%s

#### **Parameters**

**report** (STRING): Report indicating the critical state.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## monitor.fan.failed

### Severity

ALERT

### Description

This message occurs when one or more main unit fans have failed. The system remains operational. If the condition persists too long, an automatic shutdown for overtemperature might occur.

### Corrective Action

Reseat the failed fans. If the error persists, replace them.

### Syslog Message

%s

### Parameters

**report** (STRING): Report indicating the failed fans.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## monitor.fan.ok

### Severity

NOTICE

### Description

This message occurs when a previously failed main unit fan has been restarted.

### Corrective Action

(None).

### Syslog Message

%s

### Parameters

**report** (STRING): Report indicating that the failed fans have restarted.

## monitor.fan.warning

### Severity

NOTICE

### Description

This message occurs when one or more fans are in a warning state.

### Corrective Action

Replace the indicated fans to avoid overheating.

### Syslog Message

%s

## Parameters

**report** (STRING): Report indicating the warning level.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

# monitor.fru events

## monitor.fru.info.readable

### Severity

INFORMATIONAL

### Description

This message occurs when the inventory information about a field-replaceable unit (FRU) that was previously unreadable becomes readable.

### Corrective Action

(None).

### Syslog Message

The inventory information of FRU %s is readable.

## Parameters

**fru\_name** (STRING): FRU name.

## monitor.fru.info.unreadable

### Severity

ERROR

### Description

This message occurs when the inventory information of a field-replaceable unit (FRU) is not readable.

### Corrective Action

See the system hardware guide for the replacement of the FRU. At the earliest convenient time, remove and reinsert the FRU to reseal it, and then resume normal system operation. If the error persists, replace the FRU according to the hardware guide directions.

### Syslog Message

The inventory information of FRU %s is not readable.

## Parameters

**fru\_name** (STRING): FRU name.

# monitor.globalstatus events

## monitor.globalStatus.critical

**Severity**

EMERGENCY

**Description**

This message occurs when the system's global health is critical; for example, if a fan failed.

**Corrective Action**

Check the event log for other conditions that can cause this global status message. Follow the corrective actions in those messages.

**Syslog Message**

%s

**Parameters**

**problem** (STRING): System problem description.

**monitor.globalStatus.nonCritical****Severity**

ERROR

**Description**

This message occurs when the system's global health is degraded; for example, if a file system is nearly full.

**Corrective Action**

Check the event log for other conditions that can cause this global status message. Follow the corrective actions in those messages.

**Syslog Message**

%s

**Parameters**

**problem** (STRING): System problem description.

**monitor.globalStatus.nonRecoverable****Severity**

EMERGENCY

**Description**

This message occurs when the system's global health is non-recoverable, and the system is shutting down. There are currently no examples of this failure.

**Corrective Action**

Check the event log for other conditions that can cause this global status message. Follow the corrective actions in those messages.

**Syslog Message**

%s

**Parameters**

**problem** (STRING): System problem description.

**monitor.globalStatus.ok****Severity**

NOTICE

**Description**

This message occurs when the system's global health returns to normal.

**Corrective Action**

(None).

**Syslog Message**

%s

**Parameters**

**problem** (STRING): Description of the system's normal condition.

**monitor.iocard events****monitor.ioCard.degraded****Severity**

ALERT

**Description**

This message indicates that an I/O card is degraded.

**Corrective Action**

Exchange or replace the card with a known good card in the same I/O slot. If the problem persists, contact NetApp technical support for replacement assistance.

**Syslog Message**

IO card is degraded: %s

**Parameters**

**reasonText** (STRING): Slot number of the card and the reason the card is degraded.

**monitor.ioCard.ok****Severity**

NOTICE

**Description**

This message indicates that an I/O card is operating properly.

**Corrective Action**

None.

### **Syslog Message**

IO card on slot %d is OK

### **Parameters**

**slot\_number** (INT): Slot number of the card.

## **monitor.ioexpansionpower events**

### **monitor.ioexpansionPower.degraded**

#### **Severity**

NOTICE

#### **Description**

This message indicates that power on the I/O expansion module is degraded.

#### **Corrective Action**

Degraded power might be caused by bad power supplies, bad wall power, or bad components on the motherboard. If spare power supplies are available, try exchanging them to see whether the problem is resolved. Otherwise, contact NetApp technical support.

### **Syslog Message**

Power on I/O expansion module is degraded: %s

### **Parameters**

**reasonText** (STRING): Reading and state of voltage sensors on the I/O expansion module.

### **monitor.ioexpansionPower.ok**

#### **Severity**

NOTICE

#### **Description**

This messages indicates that power on the I/O expansion module is OK.

#### **Corrective Action**

No corrective action is needed.

### **Syslog Message**

Power on the I/O expansion module is OK.

### **Parameters**

(None).

## **monitor.ioexpansiontemperature events**

### **monitor.ioexpansionTemperature.cool**



**Severity**

EMERGENCY

**Description**

This warning message occurs when the I/O expansion module is too cold.

**Corrective Action**

The system cannot function in an environment that is too cold; find ways to warm the system.

**Syslog Message**

I/O expansion module is too cold: %s

**Parameters**

**describe\_toocool** (STRING): Reading and state of the temperature sensors on the I/O expansion module.

**monitor.ioexpansionTemperature.ok****Severity**

NOTICE

**Description**

This message occurs when the temperature of the I/O expansion module is normal. It can occur for the following two cases: 1) LOG\_NOTICE to show that a bad condition has reverted to normal. 2) LOG\_INFO for hourly to indicate that the temperature is OK.

**Corrective Action**

No corrective action is needed.

**Syslog Message**

Temperature of the I/O expansion module is OK.

**Parameters**

(None).

**monitor.ioexpansionTemperature.warm****Deprecated**

Deprecated as of version 9.2.0 to address the EMS ID length violation.

**Severity**

ALERT

**Description**

This warning message occurs when the I/O expansion module is too warm.

**Corrective Action**

Evaluate the environment in which the system is functioning: are air conditioning units needed or is the current air conditioning not functioning properly?

**Syslog Message**

I/O expansion module is too warm: %s

## Parameters

**describe\_toowarm** (STRING): Reading and state of the temperature sensors on the I/O expansion module.

# monitor.ioxmtemp events

## monitor.ioxmTemp.warm

### Severity

ALERT

### Description

This message occurs when the I/O module is above the recommended operating temperature.

### Corrective Action

Evaluate the environment in which the system is functioning for proper air flow and temperature control.

### Syslog Message

I/O expansion module is too warm: %s

## Parameters

**describe\_toowarm** (STRING): Reading and state of the temperature sensors on the I/O expansion module.

# monitor.mismatch events

## monitor.mismatch.hourly

### Severity

ALERT

### Description

This message occurs when the version of ONTAP® is not supported on this hardware. You cannot run a version of ONTAP produced by one vendor on hardware supplied by another. The system will shut down if you do not install a correctly branded version of ONTAP.

### Corrective Action

Install a correctly branded version of ONTAP before the timer expires.

### Syslog Message

ONTAP detected a branding mismatch. The hardware is branded by %s and ONTAP is branded by %s. There are %d hour(s) left before shutdown. Install a version of ONTAP branded by your hardware vendor.

## Parameters

**hardware** (STRING): Vendor brand associated with the hardware.

**software** (STRING): Vendor brand associated with the version of ONTAP.

**hours** (INT): Number of hours before shutdown.

## monitor.mismatch.shutdown

### Severity

EMERGENCY

### Description

This message occurs when the system shuts down because the hardware brand does not match the ONTAP® version brand.

### Corrective Action

Install a correctly branded version of ONTAP to prevent subsequent shutdowns for brand mismatch. Monitor the EMS\_monitor\_mismatch\_hourly message to determine when the next shutdown will occur because of the brand mismatch.

### Syslog Message

Shutting down because of a software branding mismatch. The hardware is branded by %s and ONTAP is branded by %s.

### Parameters

**hardware** (STRING): Vendor brand associated with the hardware.

**software** (STRING): Vendor brand associated with the version. of ONTAP.

## monitor.nvmembattery events

### monitor.nvmembattery.warninglow

### Severity

ALERT

### Description

This message occurs when the nonvolatile memory (NVMEM) battery charge is low.

### Corrective Action

Replace the NVMEM battery as soon as possible.

### Syslog Message

NVMEM battery is low on power and should be replaced as soon as possible.

### Parameters

(None).

## monitor.nvramlowbatteries events

### monitor.nvramLowBatteries

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM batteries are at a dangerously low power level.

**Corrective Action**

Replace the NVRAM batteries.

**Syslog Message**

NVRAM batteries are dangerously low.

**Parameters**

(None).

**monitor.nvramLowBatteries.notice****Severity**

NOTICE

**Description**

This message occurs when a low batteries situation exists. The timeout[123] values are duplicated as a convenience for the Syslog message, which prints out the value three times.

**Corrective Action**

Replace the NVRAM batteries.

**Syslog Message**

If the NVRAM batteries are dangerously low, the system shuts down automatically every %d hours to encourage you to replace them. If you reboot the system it will run for another %d hours before shutting down.

**Parameters**

**timeout1** (INT): Timeout value.

**timeout2** (INT): Timeout value.

**timeout3** (INT): Timeout value.

**monitor.nvramlowbattery events****monitor.nvramLowBattery****Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery is at a dangerously low power level.

**Corrective Action**

Replace the NVRAM battery.

**Syslog Message**

NVRAM battery is dangerously low.

**Parameters**

(None).

## monitor.nvramLowBattery.notice

### Severity

NOTICE

### Description

This message occurs when a low battery situation exists. The timeout[123] values are duplicated as a convenience for the Syslog message, which prints out the value three times.

### Corrective Action

Replace the NVRAM battery.

### Syslog Message

If the NVRAM battery is dangerously low, the system shuts down automatically every %d hours to encourage you to replace it. If you reboot the system it will run for another %d hours before shutting down.

### Parameters

**timeout1** (INT): Timeout value.

**timeout2** (INT): Timeout value.

**timeout3** (INT): Timeout value.

## monitor.partnercontroller events

### monitor.partnerController.notPresent

### Severity

ALERT

### Description

This message occurs when both controllers are configured as an HA (high-availability) pair, but the partner controller is not present. The system fan will run at a higher speed to prevent overheating due to the non-optimum air flow with only one controller in place.

### Corrective Action

Put the partner controller back into the HA configuration; otherwise, the system fan will run at a higher speed to prevent overheating due to the non-optimum air flow with only one controller in place.

### Syslog Message

Both controllers are configured as an HA (high-availability) pair, but the partner controller is not present. The system fan will run at a higher speed to prevent possible overheating.

### Parameters

(None).

## monitor.power events

### monitor.power.degraded

### Deprecated

Deprecated as of version 9.2. Last used in 7.3.6 on Tin and Dell one-PSU systems. It is the last user of

snmp trap 43.

**Severity**

EMERGENCY

**Description**

This message occurs when the main unit power supply is in a degraded mode.

**Corrective Action**

(Call support).

**Syslog Message**

Power supply is in degraded mode: %s

**Parameters**

**report** (STRING): Description of the problem.

**monitor.power.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the system has sufficient power (220V) to power up both controllers.

**Corrective Action**

(None).

**Syslog Message**

The power provided to the system is sufficient.

**Parameters**

(None).

**monitor.power.not.sufficient****Severity**

ALERT

**Description**

This message occurs when a storage system in a low-line (110V) power configuration does not have sufficient power to support both controllers. The built-in hardware logic prevents the controller in the bottom slot from powering up.

**Corrective Action**

A high-line (220V) power configuration is needed to power up both controllers. Power off the system, connect the power supply unit (PSU) with 220V power, and then power up the system.

**Syslog Message**

The bottom controller is not powered up due to insufficient power.

## Parameters

(None).

## monitor.power.unreadable

### Severity

ERROR

### Description

This message occurs when a power sensor in the controller module is not readable.

### Corrective Action

Shut down the system and power-cycle the controller module. If the sensor is still not readable, replace the controller module.

### Syslog Message

A power sensor %s in the controller module is not readable.

### Parameters

**sensor\_name** (STRING): Sensor name.

## monitor.raid events

### monitor.raid.brokenDisk

### Severity

ERROR

### Description

This message occurs when the monitor detects a broken disk.

### Corrective Action

Replace the disk.

### Syslog Message

%s in RAID group "%s" %s broken.

### Parameters

**type** (STRING): Type of disk.

**rgname** (STRING): Name of the RAID group.

**verb** (STRING): Description of the broken state.

### monitor.raid.reconstruct

### Severity

NOTICE

### Description

This message occurs when a RAID reconstruct is initiated.

### Corrective Action

No user action is required for the reconstruct operation to succeed. Refrain from removing disks from the affected RAID group until reconstruction is completed.

### Syslog Message

Reconstructing broken %s in RAID group "%s".

### Parameters

**type** (STRING): Type of disk.

**rgname** (STRING): RAID group name.

## monitor.raid.reconstruct.susp

### Severity

ERROR

### Description

This message occurs when a RAID reconstruction needs to be performed on a RAID group, but reconstruction is suspended or disabled.

### Corrective Action

Contact NetApp technical support for assistance with reenabling RAID reconstruction.

### Syslog Message

Reconstruction needs to be performed on RAID group "%s", but RAID reconstruction is suspended or disabled.

### Parameters

**rgname** (STRING): RAID group name.

## monitor.raiddp events

### monitor.raiddp.vol.singleDegraded

### Severity

ERROR

### Description

This message occurs as part of an hourly system check when a double parity volume is in single degraded mode. This happens when there is a disk failure in the double parity volume and the RAID subsystem cannot find a suitable replacement disk to start reconstruction.

### Corrective Action

Add a spare disk suitable to reconstruction.

### Syslog Message

%s in RAID group "%s" is broken.

### Parameters

**type** (STRING): Type of disk: data, parity or dparity disk.

**rgname** (STRING): RAID group name.



# monitor.raidtec events

## monitor.raidtec.dbIDegraded

### Severity

ERROR

### Description

This message occurs as part of an hourly system check, when a RAID-TEC aggregate is in double degraded mode. This happens when there are two disk failures in any RAID group in the RAID-TEC aggregate and the RAID subsystem cannot find suitable replacement disks to start reconstruction. One more disk failure in the same RAID group leads to the aggregate being in its worst degraded state.

### Corrective Action

Replace the failed disks, which are identifiable by lighted fault LED lights, to start reconstruction. After reconstruction is complete, the output of the "aggr show" command shows that the aggregate is back to normal state.

### Syslog Message

%s in RAID group "%s" are broken.

### Parameters

**type** (STRING): Type of disks: data, parity, dparity, or tparity disk.

**rgname** (STRING): RAID group name.

# monitor.rtchighpower events

## monitor.rtchHighPower

### Severity

ERROR

### Description

This message occurs when the real-time clock (RTC) battery voltage is above normal level.

### Corrective Action

Correct any environmental problems, such as chassis over-temperature. If the real-time clock battery voltage is still above normal 60 minutes after environmental problems have been corrected, replace it at the next available maintenance window.

### Syslog Message

Real-time clock battery voltage is above normal (%s).

### Parameters

**reading** (STRING): Battery voltage reading (mV).

# monitor.rtclopower events

## monitor.rtcLowPower

### Severity

ERROR

### Description

This message occurs when the real-time clock (RTC) battery is at a dangerously low power level.

### Corrective Action

Replace the real-time clock battery.

### Syslog Message

Real-time clock battery is dangerously low. Replace it at the next available maintenance window.

### Parameters

(None).

## monitor.rtcnormal events

### monitor.rtcNormal

### Severity

NOTICE

### Description

This message occurs when the real-time clock (RTC) battery voltage is normal. Previously occurred fault is now cleared.

### Corrective Action

(None).

### Syslog Message

Real-time clock battery voltage is normal.

### Parameters

(None).

## monitor.rtcwarnlowpower events

### monitor.rtcWarnLowPower

### Severity

ERROR

### Description

This message occurs when the real-time clock (RTC) battery is at a "warning low" voltage level.

### Corrective Action

Correct any environmental problems, such as chassis over-temperature. If the real-time clock battery voltage is still below normal 60 minutes after environmental problems have been corrected, replace it at the

next available maintenance window.

### **Syslog Message**

Real-time clock battery voltage is below normal (%s).

### **Parameters**

**reading** (STRING): Battery voltage reading (mV).

## **monitor.shelf events**

### **monitor.shelf.accessError**

#### **Severity**

ALERT

#### **Description**

This message occurs when the monitor detects that a disk shelf access error has occurred. The message is also an hourly reminder that the error has not been corrected. Enclosure Services is unable to monitor or control one or more disk shelves on a particular Fibre Channel loop or Serial SCSI Attach (SAS) domain because it cannot communicate with the enclosure services processes in those shelves.

#### **Corrective Action**

Check the corrective actions of enclosure services error messages that might precede this error in the log. Ensure that all loop or domain connections are latched and secured properly. Correct any error conditions indicated by LEDs on any of the shelves or shelf I/O modules. This might require power-cycling of shelves or replacement of modules, as is appropriate to the particular shelf type.

### **Syslog Message**

Enclosure services has detected an error in access to shelves on channel %s.

### **Parameters**

**channel** (STRING): Adapter slot or switch port that has detected the fault; for example, "7a" or "myswitch:5".

### **monitor.shelf.accessError.ok**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the monitor detects that a previously detected shelf access error has been corrected.

#### **Corrective Action**

(None).

### **Syslog Message**

Enclosure services-detected shelf access error has been corrected on channel %s.

### **Parameters**

**channel** (STRING): Adapter slot or switch port where the error was corrected.

## monitor.shelf.configError

### Severity

ALERT

### Description

This message occurs when the Enclosure Services process notifies the monitor of a disk shelf configuration error. The message is also an hourly reminder that the error has not been corrected. Configuration errors might cause Enclosure Services to be unable to monitor a disk shelf, or might indicate shelf configuration or interconnect problems; for example, too few power supplies in a disk shelf, or improper mixing of modules in a shelf or on a loop.

### Corrective Action

See the EMS message concerning disk shelf configuration errors that preceded the first instance of this message, and follow the corrective actions therein.

### Syslog Message

Enclosure services has detected an error in access to shelves or shelf configuration on channel %s.

### Parameters

**channel** (STRING): Adapter slot or switch port that has detected the fault; for example, "7a" or "myswitch:5".

## monitor.shelf.configError.ok

### Severity

NOTICE

### Description

This message occurs when the monitor detects that a previously detected shelf configuration error has been corrected.

### Corrective Action

(None).

### Syslog Message

Enclosure services-detected configuration error has been corrected on channel %s.

### Parameters

**channel** (STRING): Adapter slot or switch port where the fault existed.

## monitor.shelf.fault

### Severity

ALERT

### Description

This message occurs when the system detects a critical disk shelf fault. It is also an hourly reminder that the condition has not been corrected.

## Corrective Action

Check the event log for more specific shelf error messages, and follow the corrective actions there. Use the "storage shelf show" and "storage shelf show -errors" commands to display the current shelf status.

## Syslog Message

Critical fault reported on disk storage shelf attached to channel %s. Check fans, power supplies, disks, and temperature sensors.

## Parameters

**channel** (STRING): Adapter slot or switch port that has detected the fault; for example, "7a" or "myswitch:5".

## monitor.shelf.fault.ok

### Severity

NOTICE

### Description

This message occurs when the monitor detects that a previously detected shelf fault has been corrected.

### Corrective Action

(None).

### Syslog Message

Fault previously reported on disk storage shelf attached to channel %s has been corrected.

### Parameters

**channel** (STRING): Adapter slot or switch port where the error has been corrected; for example, "7a" or "myswitch:5".

## monitor.shelf.warning

### Severity

ERROR

### Description

This message occurs when the system detects a disk shelf warning. It is also an hourly reminder that the condition has not been corrected.

### Corrective Action

Check the messages log for more specific shelf error messages, and follow the corrective actions there.

### Syslog Message

Fault reported on disk storage shelf attached to channel %s. Check fans, power supplies, disks, and temperature sensors.

### Parameters

**channel** (STRING): Adapter slot or switch port that has detected the fault; for example, "7a" or "myswitch:5".

# monitor.shutdown.events

## monitor.shutdown.brokenDisk

### Severity

EMERGENCY

### Description

This message occurs when an automatic shutdown sequence is initiated due to a degraded RAID group that cannot be reconstructed because there are insufficient appropriate spare disks. The timeout period is specified by the "raid.timeout" option, which must be set by user to a nonzero value in order for the system to shutdown on expiry of the timeout period. If the option is not set or is set to zero, then system will not shut down due to a degraded RAID group.

### Corrective Action

To recover from the shutdown, boot the system. Ensure that there are sufficient spare disks of the appropriate type so that the degraded RAID group can be reconstructed.

### Syslog Message

%s in RAID group "%s" %s broken. Halting system now.

### Parameters

**type** (STRING): Type of disk.  
**rgname** (STRING): RAID group name.  
**verb** (STRING): Description of the broken state.

## monitor.shutdown.brokenDisk.pending

### Severity

NOTICE

### Description

This message occurs when an automatic shutdown sequence has been postponed due to RAID reconstruction.

### Corrective Action

(None).

### Syslog Message

%s in RAID group "%s" %s broken. Halting system in %d hour%s.

### Parameters

**type** (STRING): Type of disk.  
**rgname** (STRING): RAID group name.  
**verb** (STRING): Description of the broken state.  
**countdown** (INT): Number of hours remaining to shutdown.  
**plural** (STRING): Plural character.

## monitor.shutdown.cancel

**Severity**

NOTICE

**Description**

This message occurs when an automatic shutdown sequence is canceled.

**Corrective Action**

(None).

**Syslog Message**

Automatic shutdown sequence canceled.

**Parameters**

(None).

**monitor.shutdown.cancel.brokenDisk****Severity**

NOTICE

**Description**

This message occurs when an automatic shutdown sequence has been postponed due to RAID reconstruction

**Corrective Action**

(None).

**Syslog Message**

%s in RAID group "%s" %s broken. Halt delayed until %s finishes.

**Parameters**

**type** (STRING): Type of disk.

**rgname** (STRING): RAID group name.

**verb** (STRING): Description of the broken state.

**reason** (STRING): Reason that the shutdown was postponed. Possible reasons are "RAID reconstruction" and "parity recomputation".

**monitor.shutdown.cancel.nvramLowBatteries****Severity**

NOTICE

**Description**

This message occurs when an automatic shutdown sequence has been postponed due to RAID reconstruction.

**Corrective Action**

(None).

**Syslog Message**

NVRAM batteries are dangerously low. Halt delayed until %s finishes.

## Parameters

**reason** (STRING): Reason that the shutdown was postponed. Possible reasons are "RAID reconstruct" and "parity recomputation".

## monitor.shutdown.cancel.nvramLowBattery

### Severity

NOTICE

### Description

This message occurs when an automatic shutdown sequence has been postponed due to RAID reconstruction

### Corrective Action

(None).

### Syslog Message

NVRAM battery is dangerously low. Halt delayed until %s finishes.

## Parameters

**reason** (STRING): Reason that the shutdown was postponed. Possible reasons are "RAID reconstruct" and "parity recomputation".

## monitor.shutdown.chassisOverTemp

### Severity

EMERGENCY

### Description

This message occurs when the chassis temperature is too hot. This is sent just before shutdown.

### Corrective Action

Ensure that sufficient cooling air is being supplied to the chassis and that the air inlets and outlets of the unit are not blocked.

### Syslog Message

Chassis temperature is too hot: %s

## Parameters

**describe\_toohot** (STRING): Description of the condition.

## monitor.shutdown.chassisUnderTemp

### Severity

ERROR

### Description

This message occurs when the system is shutting down because chassis temperature is too cold.



**Corrective Action**

The appliance is in an environment that is too cold. Find a way to warm the environment around the appliance.

**Syslog Message**

Chassis temperature is too cold: %s

**Parameters**

**describe\_toocold** (STRING): Description of the condition.

**monitor.shutdown.emergency****Severity**

EMERGENCY

**Description**

This message occurs when ONTAP® initiates an emergency shutdown.

**Corrective Action**

Correct the condition noted in the reason field.

**Syslog Message**

Emergency shutdown: %s

**Parameters**

**reason** (STRING): Reason for the shutdown.

**monitor.shutdown.ioexpansionOverTemp****Severity**

EMERGENCY

**Description**

This message occurs when the I/O expansion module is too hot. This message is sent just before shutdown.

**Corrective Action**

The system environment is too hot; cool the environment.

**Syslog Message**

I/O expansion module is too hot: %s

**Parameters**

**describe\_toohot** (STRING): Reading and state of the temperature sensors on the I/O expansion module.

**monitor.shutdown.ioexpansionUnderTemp****Severity**

EMERGENCY

## Description

This message occurs when the I/O expansion module is too cold. This message is sent just before shutdown.

## Corrective Action

The system environment is too cold; warm the environment.

## Syslog Message

I/O expansion module is too cold: %s

## Parameters

**describe\_toocold** (STRING): Reading and state of the temperature sensors on the I/O expansion module.

## monitor.shutdown.nvramLowBatteries

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM power in a controller with multiple NVRAM batteries is dangerously low, and ONTAP® initiates a shutdown to protect user data.

### Corrective Action

Replace the controller NVRAM batteries.

### Syslog Message

Emergency shutdown: NVRAM batteries dangerously low in degraded mode. Replace the batteries immediately!

### Parameters

(None).

## monitor.shutdown.nvramLowBatteries.pending

### Severity

ALERT

### Description

This message occurs when an automatic shutdown sequence is pending because of low NVRAM batteries voltage.

### Corrective Action

Replace the controller NVRAM batteries.

### Syslog Message

NVRAM batteries are dangerously low. Halting system in %d hour%s. Replace the batteries immediately!

### Parameters

**countdown** (INT): Hours remaining until shutdown.

**plural** (STRING): Grammatical number.

## monitor.shutdown.nvramLowBattery

### Severity

EMERGENCY

### Description

This message occurs when the controller NVRAM voltage is dangerously low, and ONTAP® initiates a shutdown to protect user data.

### Corrective Action

Replace the controller NVRAM battery.

### Syslog Message

Emergency shutdown: NVRAM battery dangerously low in degraded mode. Replace the battery immediately!

### Parameters

(None).

## monitor.shutdown.nvramLowBattery.pending

### Severity

ALERT

### Description

This message occurs when an automatic shutdown sequence is pending because of low NVRAM battery voltage.

### Corrective Action

Replace the controller NVRAM battery.

### Syslog Message

NVRAM battery is dangerously low. Halting system in %d hour%s. Replace the battery immediately!

### Parameters

**countdown** (INT): Hours remaining until shutdown.

**plural** (STRING): Grammatical number.

## monitor.sparelabelcheckfailed events

### monitor.spareLabelCheckFailed

### Severity

NOTICE

### Description

This message occurs when the system detects that a hot spare disk has an invalid or a missing label. The system corrects the situation by rewriting the label.

## Corrective Action

No user action is required.

## Syslog Message

Periodic check of hot spare %s has failed. The system will correct the problem.

## Parameters

**disk\_info** (STRING): Formatted information of the disk object.  
**label\_block\_num** (INT): Label block number code.  
**label\_state** (INT): Label state code.  
**label\_read\_error** (INT): Label read error code.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk in RPM.

# monitor.temp events

## monitor.temp.unreadable

### Severity

ERROR

### Description

This message occurs when the controller module temperature is not readable. In this condition, the system does not automatically shut down if it becomes too hot for reliable operation.

### Corrective Action

Shut down the system and power-cycle the controller module. If the temperature is still not readable, replace the controller module.

### Syslog Message

The controller temperature (%s) is not readable.

### Parameters

**sensor\_name** (STRING): Sensor name.

# monitor.unknown events

## monitor.unknown.message

### Severity

ERROR

### Description

This message occurs when the status monitor receives an unexpected message type. This is an internal

error.

### Corrective Action

(None).

### Syslog Message

Internal error: status monitor received message type %d, expected type %d

### Parameters

**msgType** (INT): (None).

**expectedType** (INT): (None).

## monitor.vol events

### monitor.vol.full.inc.sav

#### Severity

ALERT

#### Description

This message occurs when one or more file systems are full, typically indicating at least 98% full. Space usage is computed based on the active file system size as well as all of the space saved by storage efficiency features. This value is computed by subtracting the value of the "Snapshot Reserve" field, and is equivalent to the "logical-used" field of the "volume show-space" command. The volume can be over 100% full due to space used or including savings by storage efficiency features or reserved by metadata. A value greater than 100% can cause the volume to become logically overallocated. See the "vol.log.overalloc.inc.sav" event message for more information.

#### Corrective Action

Create space by increasing the volume size, deleting data, or deleting Snapshot copies. To increase volume size, use the "volume size" command. To delete Snapshot copies, use the "volume snapshot delete" command.

#### Syslog Message

%s %s%s%s (UUID %s) is full (reserved or using %d%% of space including savings and %d%% of inodes).

#### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies. This event should be applicable only to volumes with logical space enforcement enabled.

**name** (STRING): Name of this object.

**app** (STRING): Application Universally Unique ID (UUID).

**vserver\_uuid** (STRING): UUID of the object's or volume's Vserver.

**object\_uuid** (STRING): UUID of the object or volume.

**percent\_full\_blocks** (INT): Used capacity of the space of the object, as a percentage.

**percent\_full\_inodes** (INT): Used capacity of inodes of the object, as a percent.

### monitor.vol.nearFull.inc.sav

#### Severity

ALERT

## Description

This message occurs when one or more file systems are nearly full, typically indicating at least 95% full. Space usage is based on the active file system size as well as all of the space saved by storage efficiency features. This value is computed by subtracting the value of the "Snapshot Reserve" field, and is equivalent to the "logical-used" field of the "volume show-space" command.

## Corrective Action

Create space by increasing the volume size, deleting data, or deleting Snapshot copies. To increase volume size, use the "volume size" command. To delete Snapshot copies, use the "volume snapshot delete" command.

## Syslog Message

%s %s%s%s (UUID %s) is nearly full (reserved or using %d%% of space including savings and %d%% of inodes).

## Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies. This event should be applicable only to volumes with logical space enforcement enabled.

**name** (STRING): Name of this object.

**app** (STRING): Application Universally Unique ID (UUID).

**vserver\_uuid** (STRING): UUID of the object's or volume's Vserver.

**object\_uuid** (STRING): UUID of the object or volume.

**percent\_full\_blocks** (INT): Used capacity of the space of the object, including space saved by the storage efficiency features, as a percentage.

**percent\_full\_inodes** (INT): Used capacity of inodes of the object, as a percentage.

# monitor.volume events

## monitor.volume.full

### Severity

ALERT

### Description

This message occurs when one or more file systems are full, typically indicating at least 98% full. The space usage is computed based on the active file system size and is computed by subtracting the value of the "Snapshot Reserve" field from the value of the "Used" field of the "volume show-space" command. Either a volume or an aggregate can be over 100% full due to space used or reserved by metadata. A value greater than 100% might cause Snapshot™ copy space to become unavailable or cause a volume to become logically overallocated. See the "vol.log.overalloc" EMS message for more information.

### Corrective Action

Create space by increasing the volume or aggregate size, or by deleting data or deleting Snapshot copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot copies, use the "volume snapshot delete" command. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot copies are deleted automatically when the aggregate is full.

### Syslog Message

%s "%s%s%s" is full (using or reserving %s%% of space and %s%% of inodes).

## Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).  
**name** (STRING): Name of this object.  
**app** (STRING): Application UUID.  
**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the object's SVM, if the object is a volume. Otherwise, this string is empty.  
**percent\_full\_blocks** (STRING): Used capacity of the space of the object, as a percentage.  
**percent\_full\_inodes** (STRING): Used capacity of inodes of the object, as a percentage.

## monitor.volume.nearlyFull

### Severity

ERROR

### Description

This message occurs when one or more file systems are nearly full, typically indicating at least 95% full. The space usage is computed based on the active file system size and is computed by subtracting the value of the "Snapshot Reserve" field from the value of the "Used" field of the "volume show-space" command.

### Corrective Action

Create space by increasing the volume or aggregate sizes, or by deleting data or deleting Snapshot® copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot® copies, use the "volume snapshot delete" command. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot® copies are deleted automatically when the aggregate is full.

### Syslog Message

%s %s%s%s is nearly full (using or reserving %s%% of space and %s%% of inodes).

## Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).  
**name** (STRING): Name of this object.  
**app** (STRING): Application UUID.  
**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the object's Vserver, if the object is a volume. Otherwise, this string is empty.  
**percent\_full\_blocks** (STRING): Used capacity of the space of the object, as a percent.  
**percent\_full\_inodes** (STRING): Used capacity of inodes of the object, as a percent.

## monitor.weeklystatsasup events

### monitor.weeklyStatsASUP.off

### Severity

ERROR

### Description

This message occurs when weekly performance AutoSupport(tm) is disabled.

### Corrective Action

Enable the weekly performance AutoSupport by entering the "options stats.autosupport.weekly.enable on"

command. You must have the correct privilege to run this command.

### **Syslog Message**

Weekly performance AutoSupport has been disabled.

### **Parameters**

(None).



# moveagent events

## moveagent.abort events

### moveagent.abort.cutover

#### Severity

INFORMATIONAL

#### Description

This event is issued during volume move when cutover is aborted because the volume took too long to replicate. The volume is unquiesced.

#### Corrective Action

The move job can be in process of retry. User can try '-force-cutover true' to force cutover.

#### Syslog Message

Cutover abort due to volume quiesce timer expiration: volume dsid %llu.

#### Parameters

**volume\_dsid** (LONGINT): The source volume data set identifier (dsid)

## moveagent.cleanup events

### moveagent.cleanup.dest

#### Severity

ERROR

#### Description

This message occurs when a volume move operation running on the partner cluster is interrupted by a disaster recovery event (such as switchover or switchback) in the MetroCluster(tm) configuration before the operation reaches the cutover phase. This leaves the temporary destination volume created on this cluster in an incomplete state. The temporary destination volume must be deleted.

#### Corrective Action

If necessary, identify which temporary destination volume must be deleted by running the "volume show -fields volume,vserver -dsid" command, and then take that volume offline by running the "volume offline" command. Delete the volume by running the "volume delete" command while it is offline.

#### Syslog Message

MetroCluster disaster recovery event requires deletion of the temporary destination volume (DSID: %llu) created by a recent volume move operation running on the partner cluster. Failure to delete the temporary destination volume might cause future switchback operations to fail.

#### Parameters

**volume\_dsid** (LONGINT): Destination volume data set ID (DSID).

## moveagent.cleanup.src

### Severity

ERROR

### Description

This message occurs when a volume move operation is interrupted by a disaster recovery event (such as switchover or switchback) in the MetroCluster(tm) configuration after the cutover phase has started but before the source volume is deleted. The source volume must be deleted.

### Corrective Action

If necessary, identify which source volume must be deleted by running the "volume show -fields volume, vservers -dsid" command, and then take the volume offline by running the "volume offline" command. Delete the volume by running the "volume delete" command while it is offline.

### Syslog Message

A MetroCluster disaster recovery event interrupted the volume move operation before the source volume was deleted. Delete the source volume (DSID: %llu). Failure to delete the source volume might cause future switchback operations to fail.

### Parameters

**volume\_dsid** (LONGINT): Source volume data set ID (DSID).

## moveagent.commitdst events

### moveagent.commitdst.fail

### Severity

ALERT

### Description

This message occurs when a 'volume move' operation finishes moving the volume to the destination but fails to update the appropriate information on the destination. The system continues to retry the operation until it is successful. In the time between the failure and a successful retry, there might be disruption to client traffic associated with the volume being moved.

### Corrective Action

Monitor the volume move job. If the problem persists, contact NetApp technical support

### Syslog Message

Unable to finish the volume move operation for DSID %llu. Failed to commit the destination volume.

### Parameters

**volume\_dsid** (LONGINT): Destination volume Data Set ID (DSID).

## moveagent.commitsrc events

### moveagent.commitsrc.fail

**Severity**

ALERT

**Description**

This message occurs when an active volume move operation cannot commit the source volume. The system continues to retry the operation until it is successful. In the time between the failure and a successful retry, there might be disruption to client traffic associated with the volume being moved.

**Corrective Action**

Monitor the volume move job. If the problem persists, contact NetApp technical support

**Syslog Message**

Unable to finish the volume move operation for DSID %llu. Failed to commit the source volume.

**Parameters**

**volume\_dsid** (LONGINT): Source volume Data Set ID (DSID).

## moveagent.cutovercomp events

### moveagent.cutovercomp.fail

**Severity**

ALERT

**Description**

This message occurs when a 'volume move' operation cannot complete cutover. The system continues to retry the operation until it is successful.

**Corrective Action**

Monitor the volume move job. If the problem persists, contact NetApp technical support

**Syslog Message**

Unable to finish the volume move operation for DSID %llu. Failed to complete cutover.

**Parameters**

**volume\_dsid** (LONGINT): Destination volume Data Set ID (DSID).

## moveagent.detected events

### moveagent.detected.on.MCC.SO

**Deprecated**

This EMS is no longer needed as the issue that caused the volume inconsistency after an MCC switchback operation is fixed.

**Severity**

ERROR

## Description

This message occurs when a volume move operation is interrupted by a disaster after the critical cutover phase has started but before the source volume is deleted. It might lead to volume inconsistencies and data access disruption on the disaster site when a MetroCluster® switchback operation is performed.

## Corrective Action

To avoid data access disruption after MetroCluster switchback, some additional cleanup steps might need to be taken on the disaster site. Please run the commands 'debug vreport show' and 'debug vreport fix' on the disaster site to perform the cleanup. Contact NetApp technical support for further assistance if needed.

## Syslog Message

A disaster event interrupted the volume move operation in the critical cutover phase before the source volume %s was deleted.

## Parameters

**volname** (STRING): Source volume name.

# moveagent.dst events

## moveagent.dst.partial.commit

### Severity

ALERT

### Description

This message occurs when the move agent detects that the attributes of the destination volume are partially updated during the commit destination phase. In such scenarios, the volume move operation may not be able to complete successfully and the client access will be blocked during this whole period.

### Corrective Action

Monitor the volume move job using the 'volume move show' command. If the problem persists, please contact NetApp technical support for recovery procedures.

### Syslog Message

Unable to finish the volume move operation as the destination volume with dsid:%llu and uuid:%s is partially committed.

### Parameters

**volume\_dsid** (LONGINT): The destination volume data set identifier (dsid)

**volume\_uuid** (STRING): The destination volume uuid.

# moveagent.exceed events

## moveagent.exceed.window

### Severity

NOTICE

### Description

This message occurs when the time window for a volume move operation exceeds the specified limit.

## Corrective Action

(None).

## Syslog Message

Cutover window was exceeded for volume with source DSID %llu in Vserver with ID %llu Total elapsed cutover time elapsed is %llums; cutover window is %llums. Time spent in committing the source volume is %llums and for destination cutover operations is %llums.

## Parameters

**volume** (LONGINT): Volume data set ID(DSID).

**vserver** (LONGINT): Vserver ID where the move takes place.

**total\_cutover\_time** (LONGINT): Total cutover time elapsed.

**cutover\_window** (LONGINT): Predetermined cutover window.

**source\_commit\_time** (LONGINT): Time for the source volume to be committed on disk.

**destination\_operation\_latency** (LONGINT): Time for commit destination and swap VLDB.

# moveagent.outofmemory events

## moveagent.outofmemory

### Severity

NOTICE

### Description

This message occurs when the move agent fails to allocate memory for the move object. The system continues to retry the operation until it is successful.

### Corrective Action

If the problem persists, it indicates a system-level problem. Contact NetApp technical support.

### Syslog Message

Unable to finish the volume move operation for DSID %llu. Failed to allocate sufficient memory.

### Parameters

**volume\_dsid** (LONGINT): Source or destination volume data set ID (DSID).

# moveagent.phase events

## moveagent.phase.queryerr

### Severity

NOTICE

### Description

This message occurs when an active volume move operation cannot access the volume move phase information.

### Corrective Action

Monitor the currently running volume move job using the 'volume move show' clustered Data ONTAP command. If the volume move completes successfully, no further action is necessary. If the volume move

fails, you can retry the volume move or contact NetApp technical support for assistance.

### **Syslog Message**

Cannot determine the vol move phase information for volume with DSID %llu.

### **Parameters**

**volume\_dsid** (LONGINT): Destination volume data set identifier(DSID).

## **moveagent.swapvldb events**

### **moveagent.swapvldb.fail**

#### **Severity**

ALERT

#### **Description**

This message occurs when a 'volume move' operation finishes moving the volume to the destination but fails to update the volume location database (VLDB). Failing to update the VLDB causes continued attempts to access the original location instead of the new location. The system continues to retry the operation until it is successful. In the time between the failure and a successful retry, there might be disruption to data access on the volume being moved.

#### **Corrective Action**

Monitor the volume move job for 30 minutes by running the "volume move show -vserver <vservname> -volume <volName%gt;" command. If the cutover phase continues to fail while trying to swap VLDB records, contact NetApp technical support.

### **Syslog Message**

Unable to finish the volume move operation for DSID %llu. Failed to complete the swap VLDB operation successfully.

### **Parameters**

**volume\_dsid** (LONGINT): Destination volume data set ID (DSID).

### **moveagent.swapvldb.sync.fail**

#### **Severity**

ERROR

#### **Description**

This message occurs when a 'volume move' operation finishes moving the volume to the destination but fails to replicate swapping of volume location database (VLDB) to the MetroCluster(tm) sync destination.

#### **Corrective Action**

Monitor the volume move job. If the problem persists, contact NetApp technical support

### **Syslog Message**

Unable to finish replicating swap VLDB operation to MetroCluster(tm) destination for the volume move operation for DSID %llu.

## Parameters

**volume\_dsid** (LONGINT): Destination volume Data Set ID (DSID).

# moveagent.switchback events

## moveagent.switchback.veto

### Severity

INFORMATIONAL

### Description

This message occurs when the volume move operation vetoes a MetroCluster(tm) switchback because the operations are mutually exclusive.

### Corrective Action

Reissue the "metrocluster switchback" command to initiate the MetroCluster switchback after the volume move operation is complete. Or use the "volume move abort" command to abort the volume move operation. Monitor the currently running volume move operation by using the "volume move show" command.

### Syslog Message

The volume move operation in progress prevents MetroCluster switchback. Monitor the currently running volume move operation by using the "volume move show" command.

## Parameters

(None).

# moveagent.switchover events

## moveagent.switchover.veto

### Severity

INFORMATIONAL

### Description

This message occurs when the volume move operation vetoes a MetroCluster(tm) switchover because the operations are mutually exclusive.

### Corrective Action

Reissue the "metrocluster switchover" command to initiate the MetroCluster switchover after the volume move operation is complete. Or use the "-override-vetoes true" option to abort the volume move. Administrators can also abort running volume move operations by using the "volume move abort" command. Monitor the currently running volume move operation by using the "volume move show" command.

### Syslog Message

The volume move operation in progress prevents MetroCluster switchover. Monitor the currently running volume move operation by using the "volume move show" command.

**Parameters**

(None).



# mpt events

## mpt.action events

### mpt.action.reset.bus

**Severity**

NOTICE

**Description**

This message occurs when the system bus is reset in the message processing technology (MPT) adapter driver.

**Corrective Action**

(None).

**Syslog Message**

Bus reset on path %s for device status %s.

**Parameters**

**path** (STRING): Device path.

**msg** (STRING): Detail about device status.

### mpt.action.reset.dev

**Severity**

NOTICE

**Description**

This message occurs when the system device is reset in the message processing technology (MPT) adapter driver.

**Corrective Action**

(None).

**Syslog Message**

Device reset on path %s for device status %s.

**Parameters**

**path** (STRING): Device path.

**msg** (STRING): Detail about device status.

## mpt.bus events

### mpt.bus.reg

**Severity**

NOTICE

**Description**

This message occurs when the message processing technology (MPT) adapter driver registration fails.

**Corrective Action**

(None).

**Syslog Message**

%s : Bus registration failed.

**Parameters**

**mpt** (STRING): MPT device name for which registration failed.

## mpt.busreg events

### mpt.busreg.phydisk.failed

**Severity**

ERROR

**Description**

This message occurs when physical disk bus registration fails in the message processing technology (MPT) adapter driver.

**Corrective Action**

Check whether the device is connected correctly.

**Syslog Message**

%s : Physical disk bus registration failed.

**Parameters**

**mpt** (STRING): MPT device name on the path that failed.

## mpt.path events

### mpt.path.failure

**Severity**

ERROR

**Description**

This message occurs when the system cannot allocate a path in the message processing technology (MPT) adapter driver.

**Corrective Action**

Check whether the device is connected correctly.

**Syslog Message**

%s : Unable to allocate path.

## Parameters

**mpt** (STRING): MPT device name on the path that failed.

# mpt.phydisk events

## mpt.phydisk.path.failed

### Severity

ERROR

### Description

This message occurs when the system cannot allocate a physical disk path in the message processing technology (MPT) adapter driver.

### Corrective Action

Check whether the device is connected correctly.

### Syslog Message

%s : Unable to allocate physical disk path.

## Parameters

**mpt** (STRING): MPT device name on the path that failed.

# mpt.tmf events

## mpt.tmf.req

### Severity

NOTICE

### Description

This message occurs when the system cannot allocate a SCSI Task Management Function (TMF) request in the message processing technology (MPT)adapter driver.

### Corrective Action

(None).

### Syslog Message

Unable to allocate dedicated TMF request for device %s.

## Parameters

**mpt** (STRING): MPT device name.

# msata events

## msata.boot events

### msata.boot.device.failed

**Severity**

ERROR

**Description**

This message occurs when Data ONTAP® module that is responsible for monitoring the health of the mSATA device determines that the associated boot device is reporting a failure condition.

**Corrective Action**

Replace the device.

**Syslog Message**

Boot device %s has become read-only.

**Parameters**

**device** (STRING): Name of the device.

### msata.boot.device.pfa

**Severity**

ERROR

**Description**

This message occurs when Data ONTAP® determines that failure is forthcoming for the associated mSATA boot device.

**Corrective Action**

Replace the device.

**Syslog Message**

Failure predicted for boot device %s.

**Parameters**

**device** (STRING): Name of the device.

## msata.unable events

### msata.unable.to.monitor

**Severity**

ERROR

**Description**

This message occurs when Data ONTAP® module that is responsible for monitoring the health of the

mSATA device cannot extract health information from the associated device.

**Corrective Action**

Replace the device. If the problem persists, contact NetApp technical support.

**Syslog Message**

Data ONTAP could not monitor device %s.

**Parameters**

**device** (STRING): Name of the device.

# msrespmon events

## msrespmon.nfs events

### msRespMon.nfs.unresponsive

#### Severity

NOTICE

#### Description

This message occurs when data IP is detected to be unresponsive. NFS service may be affected and unreachable.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s (ID: %u) hosted on a node %s (ID: %u, Uuid: %s) has been detected to be unresponsive.

#### Parameters

**dataip** (STRING): The IP address of the data IP.

**dataipID** (INT): Unique identifier of the data IP.

**nodeName** (STRING): Name of the node hosting given data IP.

**nodeID** (INT): Unique identifier of the node hosting given data IP.

**nodeUuid** (STRING): Universal unique identifier of the node hosting given data IP.

# multi events

## multi.dsk events

### multi.dsk.cksum.err

#### Severity

EMERGENCY

#### Description

This message occurs when multiple errors are detected during checksum verification, and the adapter is not taken offline or the "disk.offline\_adapter.enable" option is disabled.

#### Corrective Action

Contact Contact NetApp technical support. for assistance with troubleshooting the reason for the write verification errors.

#### Syslog Message

Checksum verification failures reported by multiple disks connected to adapter %s.

#### Parameters

**adapterName** (STRING): Name of the adapter.

# mw events

## mw.hir events

### mw.HiR.maxref.notify

#### Severity

NOTICE

#### Description

This message occurs when extent reference has reached a pre-determined threshold value. This is a temporary notification till we support OVERFLOW file - where we can support infinite reference.

#### Corrective Action

Limit the # of clone/snapshot operation.

#### Syslog Message

This message occurs when extent reference has reached a pre-determined threshold value. Limit the # of clone/snapshot operation till we support OVERFLOW file in metawaf. Aggr: %s%s%s EID: 0x%llx

#### Parameters

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**eid** (LONGINT): EID on which maxref is reached.

## mw.iron events

### mw.iron.mount.idlog.purge

#### Severity

INFORMATIONAL

#### Description

This message is used to notify idlog purge at mount after 7 consecutive unsuccessful iron runs.

#### Corrective Action

(None).

#### Syslog Message

Volume[%s] : Wafiron mount performance may be significantly impacted, because idlog purge is triggered after[%d] unsuccessful iron runs.

#### Parameters

**vol** (STRING): Name of the volume.

**generation** (INT): Iron generation number.



# nameserv events

## nameserv.conf events

### nameserv.conf.limit.reached

#### Severity

ERROR

#### Description

This message occurs when a UNIX user or group name exceeds 64 characters. This check is performed on the last node upgrade from Data ONTAP® 8.2 to 8.3. The names that exceed this limit will not get resolved in Data ONTAP 8.3.

#### Corrective Action

Replace the name with a name that is no longer than 64 characters.

#### Syslog Message

%s exceeds 64 characters for Vserver %s and configuration %s.

#### Parameters

**name** (STRING): Name of the user or group that exceeded 64 characters.

**vserver\_name** (STRING): Name of the Vserver.

**ns\_config\_name** (STRING): Name of the name service configuration (unix-user or unix-group).

## nameserv.config events

### nameserv.config.write.failed

#### Severity

ERROR

#### Description

This message occurs when a file write for name service configuration fails.

#### Corrective Action

Autocorrection is triggered for failing Vserver, configuration and node. Ensure that the root volume File System(FS) is available and has enough space and check the status using "vserver services name-service file-version show". Contact technical support if the problem persists.

#### Syslog Message

%s configuration failed for Vserver %s on node %s.

#### Parameters

**vserver\_name** (STRING): Name of the Vserver.

**ns\_config\_name** (STRING): Name of the name service configuration.

**node\_name** (STRING): Name of the node.

# nameserv.ddns.events

## nameserv.ddns.bulk.delete

### Severity

ALERT

### Description

This message occurs when a request to delete the fully qualified domain name (FQDN) of all or some of the logical interfaces of a Vserver from the DNS server fails, resulting in stale records on the DNS server. This can lead to DNS lookup failures, resulting in mount problems.

### Corrective Action

Use the (privilege: advanced) "vserver services name-service dns dynamic-update record delete" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

### Syslog Message

DDNS delete operation has failed for all or some of the logical interfaces having FQDN "%s" for Vserver "%s".

### Parameters

**FQDN** (STRING): FQDN for which the operation failed.  
**vserver\_name** (STRING): Name of the Vserver.

## nameserv.ddns.bulk.failure

### Severity

ALERT

### Description

This message occurs when a dynamic request to update the fully qualified domain name (FQDN) of all or some of the logical interfaces of a Vserver in the DNS server fails after repeated retries, resulting in stale records on the DNS server. This can lead to DNS lookup failures, resulting in mount problems.

### Corrective Action

Use the (privilege: advanced) "vserver services name-service dns dynamic-update record" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

### Syslog Message

DDNS %s operation has failed for all or some of the logical interfaces having FQDN "%s" for Vserver "%s".

### Parameters

**Operation** (STRING): Type of operation.  
**FQDN** (STRING): FQDN for which the operation failed.  
**vserver\_name** (STRING): Name of the Vserver.

## nameserv.ddns.bulk.warning

**Severity**

ERROR

**Description**

This message occurs when a dynamic request to update the fully qualified domain name (FQDN) of all or some of the logical interfaces of a Vserver on the DNS server fails, resulting in stale records on the DNS server. This can lead to DNS lookup failures, resulting in mount problems.

**Corrective Action**

Data ONTAP® retries the failed dynamic DNS update operation. Check the DNS record to ensure that it is up to date. Use the (privilege: advanced) "vserver services name-service dns dynamic-update record" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

**Syslog Message**

DDNS %s operation has failed for all or some of the logical interfaces having FQDN "%s" for Vserver "%s".

**Parameters**

**Operation** (STRING): Type of operation.

**FQDN** (STRING): FQDN for which the operation failed.

**vserver\_name** (STRING): Name of the Vserver.

**nameserv.ddns.lif.delete****Severity**

ALERT

**Description**

This message occurs when an attempt to delete a fully qualified domain name (FQDN) from the DNS server fails, resulting in a stale record on the server. The FQDN belongs to the logical interface (LIF) of the given Vserver. Retaining the stale record can lead to DNS lookup failures, resulting in mount problems.

**Corrective Action**

Use the (privilege: advanced) "vserver services name-service dns dynamic-update record delete" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

**Syslog Message**

DDNS delete operation of FQDN "%s" failed for LIF with IP %s on Vserver "%s".

**Parameters**

**FQDN** (STRING): FQDN for which the operation failed.

**IP** (STRING): IP address for which the operation failed.

**vserver\_name** (STRING): Name of the Vserver.

**nameserv.ddns.lif.failure****Severity**

ALERT

## Description

This message occurs when an attempt to update a fully qualified domain name (FQDN) on the DNS server fails after repeated retries, resulting in a stale record on the server. The FQDN belongs to the logical interface (LIF) of the given Vserver. Retaining the stale record can lead to DNS lookup failures, resulting in mount problems.

## Corrective Action

Use the (privilege: advanced) "vserver services name-service dns dynamic-update record" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

## Syslog Message

DDNS %s operation of FQDN "%s" failed for LIF with IP %s on Vserver "%s".

## Parameters

**Operation** (STRING): Type of operation.  
**FQDN** (STRING): FQDN for which the operation failed.  
**IP** (STRING): IP address for which the operation failed.  
**vserver\_name** (STRING): Name of the Vserver.

## nameserv.ddns.lif.warning

### Severity

ERROR

### Description

This message occurs when an attempt to update a fully qualified domain name (FQDN) on the DNS server fails, resulting in a stale record on the server. The FQDN belongs to the logical interface (LIF) of the given Vserver. Retaining the stale record can lead to DNS lookup failures, resulting in mount problems.

### Corrective Action

Data ONTAP® retries the failed dynamic DNS update operation. Check the DNS record to ensure that it is up to date. Use the (privilege: advanced) "vserver services name-service dns dynamic-update record" command to attempt a manual update of the record for the LIF. If this fails, check documentation for dynamic DNS updates to correct the failure or update the record directly on the DNS server.

### Syslog Message

DDNS %s operation of FQDN "%s" failed for LIF with IP %s on Vserver "%s".

### Parameters

**Operation** (STRING): Type of operation.  
**FQDN** (STRING): FQDN for which the operation failed.  
**IP** (STRING): IP address for which the operation failed.  
**vserver\_name** (STRING): Name of the Vserver.

## nameserv.netgrpbyhost events

### nameserv.netgrpbyhost.failed

**Severity**

ERROR

**Description**

This message occurs when netgroup is configured but the netgroup byhost database is not generated during an upgrade operation. Absence of the netgroup byhost database might impact performance or access during subsequent netgroup-related data access attempts.

**Corrective Action**

Use the "vserver services name-service netgroup load" command to reload the netgroup configuration file. Use the (privilege: advanced) "vserver services name-service netgroup status" command to verify that the netgroup file was successfully loaded and the netgroup byhost database successfully generated. If there are any syntax errors after the load operation, fix them and reload the netgroup file.

**Syslog Message**

Netgroup byhost database failed to generate for Vserver %s on node %s with error %s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**node\_name** (STRING): Name of the node.

**error\_string** (STRING): Error encountered during creation of netgroup byhost.

# nand events

## nand.predict events

### nand.predict.fail

#### Severity

ALERT

#### Description

This message occurs when Data ONTAP® determines that failure is forthcoming for the associated flash device on the NVRAM module.

#### Corrective Action

Replace the device to prevent future failures. Contact technical support for assistance.

#### Syslog Message

NVRAM destage Failure predicted for device: %s.

#### Parameters

**device** (STRING): Name of the device.

## nand.unable events

### nand.unable.to.monitor

#### Severity

ALERT

#### Description

This message occurs when Data ONTAP® module that is responsible for monitoring the health of the NVRAM device cannot extract health information. This failure might indicate the NVRAM device is not operating properly.

#### Corrective Action

Use the "system environment sensors show" command to verify the device's status. If the failure still exists, contact technical support for assistance.

#### Syslog Message

Unable to monitor NVRAM device: %s.

#### Parameters

**device** (STRING): Name of the device.

# nblade events

## nblade.aclcompliance events

### Nblade.AclCompliance

#### Severity

ALERT

#### Description

This message occurs when the nfs "-v4-acl-compliance" option is enabled but not functional because the SnapMirror Synchronous secondary cluster or FlexCache origin cluster is running an ONTAP version older than 9.14.1

#### Corrective Action

Upgrade the secondary cluster to ONTAP 9.14.1 or later to use the ACL compliance feature or disable the "-v4-acl-compliance" option by running the "vserver nfs modify "-v4-acl-compliance disabled" command.

#### Syslog Message

NFS "-v4-acl-compliance" option is enabled but not functional because the SnapMirror Synchronous secondary cluster or FlexCache origin cluster is running an ONTAP version older than 9.14.1

#### Parameters

(None).

## nblade.callbacktimedout events

### Nblade.CallbackTimedOut

#### Severity

ERROR

#### Description

This message occurs when the Nblade is unable to successfully complete a PORTMAP/SM call to the client within 1 seconds. The PORTMAP/SM calls are done as part of the per-aggregate lock reclaim process after a Dblade reboot. If a large number of such calls fail for a particular aggregate, then it delays the completion of the lock reclaim process for that aggregate. That delay prevents the aggregate from granting new locks to NFS clients.

#### Corrective Action

Check if the RPC program number displayed above with the correct program version is running on the client and if not start it. You can use "rpcinfo -p" on the client to check that. It is also possible that the client is unreachable due to firewall/network connectivity issues.

#### Syslog Message

SM NOTIFY: Vserver %d, Vif %d: %s program (Program number:%d Program version:%d) on client %s is not responding.

## Parameters

**virtualServer** (INT): Identifier for the Vserver associated with this operation.  
**lifid** (INT): Identifier for the logical interface associated with this operation.  
**program** (STRING): Client program not responding.  
**programNum** (INT): RPC Program Number.  
**programVersion** (INT): RPC Program Version.  
**remotelpAddress** (STRING): IP address of the client machine.

# nblade.cifs events

## Nblade.cifs.budgetAllocFailure

### Severity

ERROR

### Description

This message occurs when there is a budget allocation failure for a CIFS subsystem because other CIFS subsystems are consuming too much memory. The three displayed CIFS subsystem have consumed most of the memory.

### Corrective Action

Consider distributing the CIFS load by moving data LIFs, adding data LIFs, and redirecting workloads, or by shutting down nonessential workloads to alleviate the memory pressure. If the problem persists after addressing issues on the network or the client, contact NetApp technical support.

### Syslog Message

Memory Allocation failed for %s. The CIFS subsystem on this node has exceeded its allotment of %llu bytes of node memory. CIFS subsystems that have consumed the most memory are %s.

## Parameters

**budgetname** (STRING): Name of the CIFS subsystem for which memory allocation failed.  
**TotalCifsBudget** (LONGINT): Amount of memory available to the CIFS subsystem.  
**topConsumers** (STRING): CIFS subsystem that have consumed the most memory, and the amount of memory consumed.

## Nblade.cifs.logon.audit.fail

### Severity

ERROR

### Description

This message occurs when a CIFS logon fails when guaranteed audit is enabled and audit log generation fails due to no space remaining in either destination volumes or staging volumes.

### Corrective Action

Use the command "event log show -messagename adt.service.block" to check for destination volume space unavailability. If the message entry exists, use the command "event log show -messagename adt.service.block -fields action" and follow the corrective action. If the message entry does not exist, then staging volumes are full and no corrective action is necessary. The audit process self-heals and normal auditing continues in 10 seconds after audit logging cleanup occurs.



## Syslog Message

CIFS logon failed due to audit log generation failure for the vserver %u.

## Parameters

**vsId** (INT): Id of the vserver, on which this event occurred.

# nblade.cifsencsessaccessdenied events

## Nblade.cifsEncSessAccessDenied

### Severity

ERROR

### Description

This message occurs when a client not capable of SMB encryption tries to establish a CIFS session that requires SMB encryption.

### Corrective Action

Either ensure that the client is capable of SMB encryption or disable SMB encryption on the Vserver.

### Syslog Message

Client (IP: %s, Vserver: %d) could not establish a CIFS session because SMB encryption is required.

### Parameters

**IpAddress** (STRING): IP address of the client attempting to establish a CIFS session.

**vserverId** (INT): Identifier of the Vserver associated with this operation.

# nblade.cifsencshraccessdenied events

## Nblade.cifsEncShrAccessDenied

### Severity

ERROR

### Description

This message occurs when a client not capable of SMB encryption tries to connect to a CIFS share that requires SMB encryption.

### Corrective Action

Either ensure that the client is capable of SMB encryption or disable SMB encryption on the CIFS share.

### Syslog Message

Client (IP: %s, Vserver: %d) could not connect to CIFS share "%s" because SMB encryption is required.

### Parameters

**IpAddress** (STRING): IP address of the client from which Share Connect requests are being attempted.

**vserverId** (INT): Identifier of the Vserver associated with this operation.

**shareName** (STRING): Name of the share to which the client is connecting.

# nblade.cifsetcdactionfailed events

## Nblade.cifsEtcdActionFailed

### Severity

ERROR

### Description

This message occurs when the DMAP container encounters an error processing an etcd request, such as a connection failure.

### Corrective Action

Check the network connectivity of the etcd server to the ONNIX pod, and correct any networking faults. Scan the DMAP and etcd logs for additional information.

### Syslog Message

Received an etcd failure event, Message is %s.

### Parameters

**pMessage** (STRING): Description of the SMB configuration update failure.

# nblade.cifsetcdmaxtxnlimit events

## Nblade.cifsEtcdMaxTxnLimit

### Severity

ERROR

### Description

This message occurs when the number of Access Control Entries (ACE) present in an MSRPC request for updating share permissions exceeds the supported maximum etcd limit of 1,022 entries.

### Corrective Action

Retry the operation with less than the maximum supported number of ACEs in the MSRPC request.

### Syslog Message

Received an etcd failure event. Message is %s.

### Parameters

**pMessage** (STRING): Description of the error.

# nblade.cifslongrunpattmatch events

## Nblade.CifsLongRunPattMatch

### Severity

ERROR

## Description

This message occurs when a CIFS query directory command that includes wildcard characters exceeds the reporting time threshold for completion. Pattern-matching searches on strings including wildcard characters consume excessive CPU resources.

## Corrective Action

Check the Windows client that is requesting searches containing wildcard characters. For assistance, contact NetApp technical support.

## Syslog Message

Detected a long running CIFS pattern matching operation. SMB dialect: %s Allowed pattern matching threshold (ms): %llu Total pattern matching time taken(ms): %llu Client IP address: %s Target Vserver name: %s Target Vserver ID: %d Wildcard Pattern: %s Share name: %s Directory Path: %s

## Parameters

**dialect** (STRING): Negotiated SMB dialect.

**allowedThresholdInMs** (LONGINT): Pattern-matching threshold value.

**totalTimeTakenInMs** (LONGINT): Total time taken for pattern matching.

**clientIpAddress** (STRING): IP address of the client machine.

**vservername** (STRING): Identifier for the Vserver name associated with this operation.

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**pattern** (STRING): Wildcard pattern sent.

**shareName** (STRING): Target share name.

**dirPath** (STRING): Target directory path.

# nblade.cifsmanyauths events

## Nblade.cifsManyAuths

### Severity

ERROR

### Description

This message occurs when many authentication negotiations occur simultaneously. There are 256 new session requests from this client that are not yet complete.

### Corrective Action

Investigate why the client is creating 256 or more new connection requests. It might be necessary to contact the vendor of the client or of the application in order to determine why this is occurring.

### Syslog Message

Many simultaneous new CIFS connections are occurring on Vserver ID %u from IP address %s object type is %s with UUID %s.

### Parameters

**vsId** (INT): ID of the Vserver on which this event occurred.

**remoteIpAddress** (STRING): IP address of the client machine.

**object\_type** (STRING): Type of the resource object under notification.

**object\_uuid** (STRING): UUID of the Vserver.

# nblade.cifsmaxopensamefile events

## Nblade.cifsMaxOpenSameFile

### Severity

ERROR

### Description

This message occurs when the number of times a user tries to open a file over a TCP connection is exceeded. This request and any further requests to open the same file by the user are denied until some open instances of the file are closed. This typically indicates abnormal application behavior.

### Corrective Action

Inspect the application running on the client using this TCP connection. The client might be operating incorrectly due to the application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set using the (privilege: advanced) "cifs option modify -max-opens-same-file-per-tree" command.

### Syslog Message

Received too many open file requests for the same file by one user on a connection: clientIP:port %s:%d, file "%s" on share "%s", vserver: "%s". Object type is: %s with UUID: %s.

### Parameters

**IpAddress** (STRING): IP address of the client from which open requests are denied.  
**port** (INT): Port number of the client from which the connection has been established.  
**filePath** (STRING): Path of the file that cannot be opened.  
**shareName** (STRING): Name of the share that holds the file.  
**vserverName** (STRING): Vserver associated with this operation.  
**object\_type** (STRING): Type of resource object.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.cifsmaxopensamefilenotice events

## Nblade.cifsMaxOpenSameFileNotice

### Severity

NOTICE

### Description

This message occurs when the number of open instances of a particular file under a CIFS tree nears the configuration limit. Upon reaching this limit, any further requests to open the same file by the user are denied until some open instances of the file are closed. This typically indicates abnormal application behavior.

### Corrective Action

Inspect the application running on the client using this TCP connection. The client might be operating incorrectly due to the application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set by using the (privilege: advanced) "cifs option modify -max-opens-same-file-per-tree" command.

## Syslog Message

Received %d open file requests, nearing the configured limit of %d, for the same file by one user on a connection: clientIP:port %s:%d, file: "%s" on share: "%s", vserver: "%s".

## Parameters

**currentCount** (INT): Current count of the open instances on the file.  
**maxValue** (INT): Maximum allowed open instances on a file.  
**IpAddress** (STRING): IP address of the client requesting the file open.  
**port** (INT): Port number of the client from which the connection has been established.  
**filePath** (STRING): Path of the file being opened.  
**shareName** (STRING): Name of the share that holds the file.  
**vserverName** (STRING): Vserver associated with this operation.

# nblade.cifsmaxsessperusrconn events

## Nblade.cifsMaxSessPerUsrConn

### Severity

ERROR

### Description

This message occurs when the number of sessions allowed per user over a TCP connection is exceeded. This request and any further session establishment requests are denied until some sessions are released. This is typically caused by a faulty client or application.

### Corrective Action

Inspect all applications running on the client, and terminate any that are not operating properly. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set using the (privilege: advanced) "cifs option modify -max-same-user-sessions -per-connection" command.

## Syslog Message

Received too many session requests from the same user on one TCP connection: clientIP:port %s:%d, user "%s", vserver: "%s". Object type is: %s with UUID: %s.

## Parameters

**IpAddress** (STRING): IP address of the client requesting a session.  
**port** (INT): Port number of the client from which the TCP connection has been established.  
**userName** (STRING): User that has exceeded session requests.  
**vserverName** (STRING): Vserver associated with this operation.  
**object\_type** (STRING): Type of resource object.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.cifsmaxsessperusrconnnotice events

## Nblade.cifsMaxSessPerUsrConnNotice

### Severity

NOTICE

## Description

This message occurs when the number of sessions allowed per user over a TCP connection nears the configured limit. Upon reaching this limit, any further session establishment requests are denied until some sessions are released. This is typically caused by a faulty client or application.

## Corrective Action

Inspect all applications running on the client, and terminate any that are not operating properly. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set by using the (privilege: advanced) "cifs option modify -max-same-user -sessions-per-connection" command.

## Syslog Message

Received %d session requests, nearing the configured limit of %d, from the same user on one TCP connection: clientIP:port %s:%d, user: "%s", vserver: "%s".

## Parameters

**currentCount** (INT): Current count of the sessions.

**maxValue** (INT): Maximum allowed sessions per user over a TCP connection.

**IpAddress** (STRING): IP address of the client requesting a session.

**port** (INT): Port number of the client from which the TCP connection has been established.

**userName** (STRING): User that is nearing the session requests limit.

**vserverName** (STRING): Vserver associated with this operation.

# nblade.cifsmaxwatchespertree events

## Nblade.cifsMaxWatchesPerTree

### Severity

ERROR

### Description

This message occurs when the total number of directory watch (Change Notify) requests exceed the per-tree limit. This request and any further directory watch requests are denied. This typically indicates abnormal client behavior.

### Corrective Action

Inspect the application using the connection and also monitor the other applications on the client. The client might be operating incorrectly due to a faulty application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which can be set using the (privilege: advanced) "cifs option modify -max-watches-set-per-tree" command.

### Syslog Message

Received too many Change Notify requests on one tree: clientIP:port %s:%d, directory "%s" on share "%s", vserver: "%s".

### Parameters

**IpAddress** (STRING): IP address of the client requesting Change Notify functionality.

**port** (INT): Port number of the client from which the TCP connection has been established.

**dirPath** (STRING): Path of the directory on which Change Notify is requested.

**shareName** (STRING): Name of the share that holds the directory.

**vserverName** (STRING): Vserver associated with this operation.

# nblade.cifsmaxwatchespertreenotice events

## Nblade.cifsMaxWatchesPerTreeNotice

### Severity

NOTICE

### Description

This message occurs when the total number of directory watch (Change Notify) requests nears the configured limit. Upon reaching this limit, any further directory watch requests are denied. This typically indicates abnormal client behavior.

### Corrective Action

Inspect the application using the connection, and also monitor the other applications on the client. The client might be operating incorrectly due to a faulty application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which can be set by using the (privilege: advanced) "cifs option modify -max-watches-set-per-tree" command.

### Syslog Message

Received %d Change Notify requests, nearing the configured limit of %d, on one tree: clientIP:port %s:%d, directory: "%s" on share: "%s", vserver: "%s".

### Parameters

**currentCount** (INT): Current count of the Change Notify on the tree.

**maxValue** (INT): Maximum allowed Change Notify on the tree.

**IpAddress** (STRING): IP address of the client requesting Change Notify functionality.

**port** (INT): Port number of the client from which the TCP connection has been established.

**dirPath** (STRING): Path of the directory on which Change Notify is requested.

**shareName** (STRING): Name of the share that holds the directory.

**vserverName** (STRING): Vserver associated with this operation.

# nblade.cifsmemexceeded events

## Nblade.cifsMemExceeded

### Severity

ERROR

### Description

This message occurs when the CIFS subsystem, which operates all the CIFS servers exposed through a given node, has requested more memory than it is allowed in this node. The subsystem is restricted to a percentage of total memory in the node, so the node itself need not be low on memory for this to occur. CIFS memory is not allocated in advance and so CIFS memory requests compete with other system processes. This message might not appear if a memory request from CIFS cannot be honored because the node is already low on memory. This might result in CIFS client operation failures.

### Corrective Action

Compare the number of memory allocation failures with the previous occurrence of this event to understand the load on the system. Further analyze the memory allocated to the CIFS subsystem on this node by viewing the nblade\_cifs object instance for this node through the "statistics show" interface. Based on that data, consider distributing CIFS load by moving data LIFs, adding data LIFs, and redirecting workloads, or

by shutting down non-essential workloads to alleviate the memory pressure.

### Syslog Message

The CIFS subsystem on this node has exceeded its allotment of %llu bytes of node memory with currently %llu memory allocation failures since boot time. This might result in unexpected CIFS application failures.

### Parameters

**availableMemoryInBytes** (LONGINT): Amount of memory available to the CIFS subsystem.

**memAllocFailures** (LONGINT): Number of memory allocation failures in the CIFS subsystem since boot time.

## nblade.cifsnbnameconflict events

### Nblade.cifsNbNameConflict

#### Severity

ERROR

#### Description

This message occurs when the NetBIOS Name Service receives a negative response from a remote machine for a name registration request. This is typically caused by a conflict in the NetBIOS name or an alias. As a result, clients might not be able to access data or connect to the right data-serving node in the cluster.

#### Corrective Action

1. If applicable, delete the duplicate NetBIOS alias by using the "vserver cifs delete -aliases alias -vserver vserver" command. Alternatively, rename a NetBIOS alias by deleting the duplicate name and adding an alias with a new name by using the "vserver cifs create -aliases alias -vserver vserver" command.
2. If there are no aliases configured and there is a conflict in the NetBIOS name, then rename the CIFS server by using the "vserver cifs delete -vserver vserver" and "vserver cifs create -cifs-server netbiosname" commands. Note that deleting a CIFS server can cause data to be inaccessible.
3. Remove or rename the NetBIOS name on the remote machine.

#### Syslog Message

The NetBIOS Name Service received a negative name registration response. The name %s is owned by a remote machine. The IP address being registered is %s. Object type is: %s with UUID: %s.

#### Parameters

**nbName** (STRING): NetBIOS name or alias being registered by the NetBIOS Name Service.

**IpAddress** (STRING): IP address being registered by the NetBIOS Name Service.

**object\_type** (STRING): Type of resource object.

**object\_uuid** (STRING): UUID of the resource object.

## nblade.cifsnonaesnibytesproc events

### Nblade.cifsNonAesNiBytesProc

#### Severity

ALERT



## Description

This message occurs when a system with CPUs that do not support the AES-NI instruction set is performing SMB3 encryption or signing of a large amount of data. These cryptographic operations can degrade performance severely for both SMB and other system operations.

## Corrective Action

Avoid encrypted or signed SMB3 traffic on systems that do not support the AES-NI instruction set.

## Syslog Message

Excessive SMB3 cryptographic processing load (AES-NI instruction set is not supported).

## Parameters

(None).

# nblade.cifsnonunicoderequest events

## Nblade.cifsNonUnicodeRequest

### Severity

ERROR

### Description

This message occurs when a client attempts to send a non-unicode request to a Vserver. The request is rejected.

### Corrective Action

Check the client to verify that it supports unicode characters and negotiates unicode support with the Vserver.

### Syslog Message

Vserver ID %u rejected non-unicode request from client IP: %s.

### Parameters

**vserverId** (INT): Identifier for the Vserver that received the non-unicode request.

**clientIp** (STRING): IP address of the client sending the non-unicode request.

# nblade.cifsnoprivshare events

## Nblade.cifsNoPrivShare

### Severity

EMERGENCY

### Description

This message occurs when a client attempts to connect to a nonexistent ONTAP\_ADMIN\$ share.

### Corrective Action

Ensure that the vsan is enabled for the mentioned Vserver ID. Enabling vsan on a Vserver causes the ONTAP\_ADMIN\$ share to be created for the Vserver automatically.

## Syslog Message

Vserver ID: %d, user name: %s, client ip: %s, Object type is: %s with UUID: %s.

## Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**userName** (STRING): User name of the client attempting to access the nonexistent ONTAP\_ADMIN\$ share.

**clientIp** (STRING): IP address of the client attempting to access the nonexistent ONTAP\_ADMIN\$ share.

**object\_type** (STRING): Type of resource object.

**object\_uuid** (STRING): UUID of the resource object.

# nblade.cifsoperationtimedout events

## Nblade.CifsOperationTimedOut

### Severity

ERROR

### Description

This message occurs when a CIFS operation times out because it could not be processed completely within 40 seconds.

### Corrective Action

Check the QoS policy configuration for the target volume. If there are no QoS policies set, check the load on the target volume and check IP connectivity over the cluster interfaces between the two nodes. For further assistance, contact NetApp technical support and indicate that you need assistance with a potential internal cluster performance problem.

## Syslog Message

Detected a timed out CIFS operation. SMB command for this operation: %s, Number of times this command was suspended: %llu, Number of times this command was restarted: %llu, Last CSM error during this operation: %s, Remote blade UUID: %s, Is QoS enabled: %s, Last nBlade error during this operation: %s, Client IP address: %s, Local IP address: %s, Target Vserver ID: %d, Target disk's DSID: %d, Target Vserver Name: %s

## Parameters

**commandName** (STRING): CIFS command for this operation.

**suspensionCnt** (LONGINT): Number of times this command has been suspended.

**cmdRestartCnt** (LONGINT): Number of times this command has been restarted.

**lastCsmError** (STRING): Last CSM error.

**remoteBladeID** (STRING): Remote blade ID for this CSM session.

**isQosEnabled** (STRING): Whether QoS is enabled.

**lastSpinNpError** (STRING): Last internal network component(nBlade) error.

**clientIpAddress** (STRING): IP address of the client machine.

**localIpAddress** (STRING): IP address of the local interface serving the protocol operation.

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**dsId** (INT): Data set ID (DSID) of the target volume.

**vserverName** (STRING): Vserver name associated with this operation.

# nblade.cifssametreepersess events

## Nblade.cifsSameTreePerSess

### Severity

ERROR

### Description

This message occurs when the number of connections to the same share allowed per session is exceeded. This request and any further requests to establish another connection to that share are denied until some connections to the share are released. This typically indicates abnormal client behavior.

### Corrective Action

Inspect the application using the connection and also monitor the other applications on the client. The client might be operating incorrectly due to a faulty application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set using (privilege: advanced) "cifs option modify -max-same-tree-connect-per-session" command.

### Syslog Message

Received too many connection requests to the same share from one session: clientIP:port %s:%d, share "%s", user "%s", vserver "%s".

### Parameters

**IpAddress** (STRING): IP address of the client from which multiple share connections have been attempted.  
**port** (INT): Port number of the client from which the TCP connection has been established.  
**shareName** (STRING): Name of the share to which the client is attempting connection.  
**userName** (STRING): User that is attempting to connect to the share.  
**vserverName** (STRING): Vserver associated with this operation.

## nblade.cifssametreepersessnotice events

### Nblade.cifsSameTreePerSessNotice

### Severity

NOTICE

### Description

This message occurs when the number of connections to the same share allowed per session nears the configured limit. Upon reaching this limit, any further requests to establish another connection to that share are denied until some connections to the share are released. This typically indicates abnormal client behavior.

### Corrective Action

Inspect the application using the connection, and also monitor the other applications on the client. The client might be operating incorrectly due to a faulty application running on it. Rebooting the client might also be helpful. In some cases, clients are operating as expected but require a higher threshold, which you can set by using the (privilege: advanced) "cifs option modify -max-same-tree-connect-per-session" command.

### Syslog Message

Received %d connection requests to the same share, nearing the configured limit of %d, from one session: clientIP:port %s:%d, share: "%s", user: "%s", vserver: "%s".

## Parameters

**currentCount** (INT): Current count of the connections on the session.  
**maxValue** (INT): Maximum allowed connections on the session.  
**IpAddress** (STRING): IP address of the client from which multiple share connections have been attempted.  
**port** (INT): Port number of the client from which the TCP connection has been established.  
**shareName** (STRING): Name of the share to which the client is attempting connection.  
**userName** (STRING): User that is attempting to connect to the share.  
**vserverName** (STRING): Vserver associated with this operation.

# nblade.cifsshrcconnectfailed events

## Nblade.cifsShrConnectFailed

### Severity

ERROR

### Description

This message occurs when the network blade cannot connect to the path specified in the CIFS share.

### Corrective Action

Verify that the CIFS share path is valid and that the volume it belongs to is accessible.

### Syslog Message

Vserver name: %s, Vserver ID: %d, error description: %s, error code: %d, share name: %s, Windows user name: %s, UNIX user name: %s, server IP address: %s, client IP address: %s.

## Parameters

**vserverName** (STRING): Name for the Vserver associated with this operation.  
**vserverId** (INT): Identifier for the Vserver associated with this operation.  
**errorDescription** (STRING): Description of the error for the CIFS share connection failure.  
**errorCode** (INT): Error code that describes the reason for the CIFS share connection failure.  
**shareName** (STRING): Name of the share to which the connection attempt failed.  
**domainWinUserName** (STRING): Windows domain\user that is attempting to connect to the share.  
**unixUserName** (STRING): UNIX user name.  
**serverIp** (STRING): IP address of the server.  
**clientIp** (STRING): IP address of the client.

# nblade.cifswitnessfonotify events

## Nblade.cifsWitnessFONotify

### Severity

INFORMATIONAL

### Description

This message occurs when the Witness system has requested its CIFS clients to move their Continuously Available connections to the local node, due to a failure of the partner node.

### Corrective Action

(None).

## Syslog Message

The Witness service received a failure notification for the partner node. Notification of %d CIFS clients to move their Continuously Available connections to this node took %d milliseconds to complete.

## Parameters

**numNotificationsSent** (INT): Number of clients notified.

**timeTakenInMsecs** (INT): Number of milliseconds taken to notify the clients.

# nblade.css events

## nblade.css.s3.AccessDenied

### Severity

ERROR

### Description

This message occurs when ONTAP® object-store-server denies access to a user.

### Corrective Action

Verify that the user has permission to access the resource. Check access and bucket policy configurations using these commands: 'vserver object-store-server policy show', 'vserver object-store-server policy show-statements', 'vserver object-store-server bucket policy show'.

## Syslog Message

Access is denied for user '%s' (Vserver %d), from client IP %s accessing resource '%s'.

## Parameters

**userName** (STRING): Name of the user.

**vserverId** (INT): ID of the Vserver.

**ipAddress** (STRING): IP address of the client.

**resource** (STRING): Resource that the user is trying to access.

# nblade.cvo events

## nblade.cvo.remote.access

### Severity

NOTICE

### Description

This message occurs when a remote volume has been accessed in the Cloud Volumes ONTAP® environment.

### Corrective Action

(None).

## Syslog Message

Detected remote volume access (Vserver: %s) from client IP %s to LIF IP %s. Volume: %s.

## Parameters

**vserver** (STRING): Vserver name.  
**client\_ipaddr** (STRING): IP address of the client.  
**lif\_ipaddr** (STRING): IP address of the LIF.  
**volume** (STRING): Name of the volume.

# nblade.dbladenoreponse events

## Nblade.dBladeNoResponse.CIFS

### Severity

ERROR

### Description

This message occurs when a CIFS network data protocol operation times out because the requesting node cannot communicate with the data-serving node to complete the file operation. If this condition persists, there is likely a problem with the communication between the node that generated the event and the node where the data is located in the cluster.

### Corrective Action

Identify the node where the data is located by using the (privilege: advanced) "system node show" command and the UUID of the node. Verify network connectivity between the node that generated the event and the node where the data is located by using the "cluster ping-cluster" and "network ping" commands.

### Syslog Message

File operation timed out because there was no response from the data-serving node. Node UUID: %s, file operation protocol: %s.

## Parameters

**DBladeUuid** (STRING): Universally Unique Identifier (UUID) of the node where this operation is trying to access.  
**protocol** (STRING): Network data protocol that timed out while attempting the file operation.

## Nblade.dBladeNoResponse.NFS

### Severity

ERROR

### Description

This message occurs when an NFS network data protocol operation times out because the requesting node cannot communicate with the data-serving node to complete the file operation. If this condition persists, there is likely to be a problem with the communication between the node that generated the event and the node where the data is located in the cluster.

### Corrective Action

Identify the node where the data is located by using the (privilege: advanced) "system node show" command and the UUID of the node. Verify network connectivity between the node that generated the event and the node where the data is located by using the "cluster ping-cluster" and the "network ping" commands.

## Syslog Message

File operation timed out because there was no response from the data-serving node. Node UUID: %s, file operation protocol: %s, client IP address: %s, Opcode/Procedure: %s(%d).

## Parameters

**DBladeUuid** (STRING): Universally unique identifier (UUID) of the node where this operation is trying to access.

**protocol** (STRING): Network data protocol that timed out while attempting the file operation.

**clientip** (STRING): IP address of client that attempted the file operation.

**nfsOpcode** (STRING): Opcode corresponding to file operation.

**fileoperation** (INT): File operation during which the timeout happened.

# nblade.didnotinitialize events

## Nblade.DidNotInitialize

### Severity

EMERGENCY

### Description

The appliance's network process failed during initialization due to an internal error (likely there is not enough system memory available to start the necessary kernel threads). Because the network process is not running, file access to volumes on local aggregates is not currently available, and cluster and management virtual interfaces for this appliance are also not currently available.

### Corrective Action

Initiate a storage failover (takeover) of any aggregates on this appliance, and perform a "sendhome" operation on these aggregates after this appliance is operational again. Also, verify that data virtual interfaces have properly migrated to another appliance. It is possible that a reboot or a reversion to a known compatible software version could also correct this problem.

## Syslog Message

The Network Blade kernel module has failed during initialization.

## Parameters

(None).

# nblade.ecv events

## nblade.ecv.not.ready

### Severity

NOTICE

### Description

This message occurs when ONTAP® software uses a default Effective Cluster Version (ECV) before the actual ECV has been conclusively determined. Some features might not work as expected until ONTAP can select the correct ECV.

## Corrective Action

(None).

## Syslog Message

Default ECV %s has been used. The ECV value '%d' does not reflect the cluster's true ECV status.

## Parameters

**ecv\_variable\_name** (STRING): Name of the ECV value being used.

**ecv\_variable\_value** (INT): ECV value being used.

# nblade.execsoverlimit events

## nblade.execsOverLimit

### Severity

NOTICE

### Description

This message occurs when the number of active requests to a connection exceeds the allowed in-flight request threshold, which results in request throttling. Client connection performance might be degraded.

### Corrective Action

Check the RPC slot setting on the client. Refer to TR-4067 "NFS Best Practice and Implementation Guide" for the recommended value and how to set it. For additional guidance, search the NetApp knowledgebase on the support web site for "nblade.execsOverLimit".

### Syslog Message

The number of in-flight requests from client IP "[%s]:%d" to destination LIF "%s" (Vserver "%d") is greater than the maximum allowed (%d).

### Parameters

**clientIpAddress** (STRING): IP address of the client.

**clientPort** (INT): Network port of the client.

**lifIpAddress** (STRING): IP address of the network interface (LIF).

**vserverId** (INT): ID of the Vserver.

**execsLimit** (INT): Maximum number of in-flight requests allowed per connection.

# nblade.exportaccesschkfailed events

## Nblade.exportAccessChkFailed

### Severity

ERROR

### Description

This message occurs when the network blade denies access to an exported share when a connection attempt violates an export policy.



### Corrective Action

Ensure that the export policy rule for this vserver is defined properly.

### Syslog Message

(None).

### Parameters

**vserverId** (INT): Identifier for the vserver associated with this operation.  
**errorCode** (INT): Error code that describes the export access check failure.  
**shareName** (STRING): Name of the share that the user attempted to access.

## nblade.exportaccessindeterm events

### Nblade.exportAccessIndeterm

#### Severity

ERROR

#### Description

This message occurs when client access cannot be evaluated because of an error while matching the client against export rules. Client access might be disrupted due to unresponsive or misconfigured external servers, problems with external server configuration on Data ONTAP®, or internal errors while processing the rules.

#### Corrective Action

Verify that external name services (such as DNS, NIS, LDAP, Kerberos, or AD) are responsive to queries from Data ONTAP®. Use the "vserver show -fields id" command to get the Vserver name corresponding to Vserver ID, and use the "vol show -fields msid" command to get the volume name corresponding to the volume MSID. Use the "vserver export-policy check-access" command to troubleshoot it further, or contact NetApp technical support for further assistance in reviewing the name services logs.

#### Syslog Message

Access to volume with MSID %llu in Vserver %u could not be evaluated for client "%s" as per export ruleset %llu.

#### Parameters

**msid** (LONGINT): MSID of the volume on which access was attempted.  
**vserverId** (INT): Identifier for the Vserver associated with this access.  
**clientAddr** (STRING): IP address of the client attempting access.  
**rulesetID** (LONGINT): Export rule set ID.

## nblade.fcvoidisconnected events

### Nblade.fcVoIDisconnected

#### Severity

EMERGENCY

#### Description

This message occurs when the NFS client attempts to access a FlexCache® volume, but the origin of the

FlexCache volume cannot be accessed due to poor network connectivity. As a result, NFS clients will lose access to the FlexCache volume.

### Corrective Action

Verify that the FlexCache volume can connect with its origin. If the FlexCache volume and origin volume are on the same cluster, then check the node health by using the "system node show" command. If the FlexCache volume and origin volume are on different clusters, then ensure that the intercluster LIFs on both clusters can connect by using the "network ping" commands to check the connectivity on both the clusters.

### Syslog Message

Attempt to access FlexCache volume with MSID %u on Vserver ID %u failed because FlexCache origin volume with MSID %u is not reachable.

### Parameters

**cacheMsid** (LONGINT): MSID of the FlexCache volume.

**cacheVserverId** (INT): Identifier for the FlexCache Vserver associated with this operation.

**originMsid** (LONGINT): MSID of the origin of the FlexCache volume.

## nblade.fileopenlimitexceeded events

### Nblade.fileOpenLimitExceeded

#### Severity

NOTICE

#### Description

This message occurs when the number of open files on a single SMB connection or session exceeds the system limit.

#### Corrective Action

(None).

#### Syslog Message

Number of open files has reached the maximum limit of (%d).

#### Parameters

**filesopen** (INT): Number of files that are open.

## nblade.flexcachecaaccess events

### Nblade.flexcacheCAAccess

#### Severity

ERROR

#### Description

This message occurs when a client tries to access a FlexCache® volume via SMB3 using a continuously available share but is denied because FlexCache does not support persistent handle capability.

## Corrective Action

Either remove the continuously available property from the CIFS share or create a new share without this property to access the FlexCache volume over SMB.

## Syslog Message

Attempt to access FlexCache volume with MSID %u via SMB3 over a continuously available share is denied.

## Parameters

**msid** (INT): MSID of the FlexCache volume being accessed.

# nblade.flexcachevolumeaccess events

## Nblade.flexCacheVolumeAccess

### Severity

ERROR

### Description

This message occurs when a client tries to access a FlexCache volume through an unsupported protocol. Currently only NFSv3,NFSv4,NFSv4.1 and SMB clients are supported by FlexCache.

### Corrective Action

Make sure a FlexCache-supported protocol is specified for all client mount operations through which a FlexCache volume can be accessed. If using FlexCache, specify 'nfsvers=3' or 'nfsvers=4' or 'nfsvers=4.1' as an option for the 'mount' command on all clients accessing a FlexCache volume.

### Syslog Message

Attempt to access FlexCache volume %s in SVM %s via %s denied for client %s.

### Parameters

**volName** (STRING): Name of the FlexCache volume being accessed.

**vserverName** (STRING): Name of the vserver associated with this operation.

**protocol** (STRING): Client protocol.

**clientIpAddress** (STRING): IP address of the client machine.

# nblade.flexgroupstatefulprotocolaccess events

## Nblade.flexgroupStatefulProtocolAccess

### Severity

ERROR

### Description

This message occurs when the system detects an attempt to use Network Lock Manager (NLM) shared locks on a FlexGroup volume. FlexGroup volumes do not support NLM shared locks.

### Corrective Action

FlexGroup volumes do not support shared locks. Consider using FlexGroup volumes that support stateful protocols such as NFSv4.0,NFSv4.1,SMB1 or SMB2 in order to use shared locks.

## Syslog Message

Attempt to access FlexGroup volumes with Master Data Set ID "%u" with "%s" is denied.

## Parameters

**msid** (INT): Master Data Set ID of the FlexGroup volume being accessed.

**protocol** (STRING): Client protocol being used.

# nblade.fpolicy events

## nblade.fpolicy.extn.failed

### Severity

EMERGENCY

### Description

This message occurs when an I/O operation directed at a FlexCache® volume fails because the file extension include and exclude list of one of the policies is too long to send to the origin of the FlexCache volume.

### Corrective Action

Decrease the number of file extensions configured in the include and exclude list in the FPolicy scope by using the "fpolicy policy scope modify" command.

## Syslog Message

FPolicy processing failed on volume MSID "%lu", for Vserver ID "%lu", policy ID "%lu", as the combined length for the file extension include and exclude list is greater than %llu. Decrease the number of file extensions configured.

## Parameters

**msid** (LONGINT): MSID of FlexCache volume.

**vserverid** (LONGINT): ID of the Vserver.

**policyid** (LONGINT): ID of the FPolicy.

**maximumListLengthAllowed** (LONGINT): Maximum length that can be sent to the origin volume.

## Nblade.fpolicy.flow.control

### Severity

ERROR

### Description

This message occurs when FPolicy flow control is engaged due to large number of file operations being monitored by FPolicy. If this condition persists, there might be an adverse effect on client operation latency.

### Corrective Action

Check the FPolicy configuration to make sure that only intended file operations are monitored and appropriate filters are applied.

## Syslog Message

FPolicy flow control is engaged as the monitored file operations rate crossed the sustainable limit[%d] for this platform.

## Parameters

**flowControlLimit** (INT): Maximum number of file operations per second that can be monitored by FPolicy.

## nblade.fpolicy.policy.skip

### Severity

EMERGENCY

### Description

This message occurs when a non-mandatory FPolicy(tm) policy check is skipped while processing an I/O on the FlexCache® volume because the file extension include and exclude list of the policy is too long to send to the origin volume of the FlexCache volume. The error does not fail the operation and other policies, if any, are matched.

### Corrective Action

Decrease the number of file extensions configured in the include and exclude list in the FPolicy scope by using the "fpolicy policy scope modify" command.

### Syslog Message

FPolicy check is skipped for a non-mandatory policy ID "%lu" on volume MSID "%lu" for Vserver ID "%lu" because the combined length for the file extension include and exclude list is greater than %llu.

## Parameters

**policyid** (LONGINT): ID of the FPolicy policy.

**msid** (LONGINT): MSID of the FlexCache volume.

**vserverid** (LONGINT): ID of the Vserver.

**maximumListLengthAllowed** (LONGINT): Maximum length that can be sent to the origin volume.

## nblade.fpolpassthruconnect events

### Nblade.fpolPassthruConnect

### Severity

INFORMATIONAL

### Description

This message occurs when a pass-through read channel is established between the node and the FPolicy server.

### Corrective Action

(None).

### Syslog Message

Pass-through read channel with the FPolicy server was established. Vserver ID [%d], FPolicy server IP address [%s].

## Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**fpserverIp** (STRING): IP address of the FPolicy server.

# nblade.fpolpassthruDisconn events

## Nblade.fpolPassthruDisconn

### Severity

INFORMATIONAL

### Description

This message occurs when the pass-through read channel between the node and the FPolicy server is disconnected.

### Corrective Action

(None).

### Syslog Message

Pass-through read channel with the FPolicy server is disconnected. Vserver ID [%d], FPolicy server IP address [%s], Disconnect reason [%s].

### Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**fpserverIp** (STRING): IP address of the FPolicy server.

**disconnReason** (STRING): Reason for disconnection.

# nblade.fpolpassthruonsmb1 events

## Nblade.fpolPassthruOnSMB1

### Severity

ERROR

### Description

This message occurs during a client attempt to establish an FPolicy pass-through read channel over SMB1, which is not supported.

### Corrective Action

Verify that both the client and Data ONTAP® support and are configured for SMB2 or later.

### Syslog Message

For Vserver "%s", the FPolicy pass-through read channel request coming from the client "%s" is rejected because it is not supported for SMB1.

### Parameters

**vserverName** (STRING): Vserver associated with this operation.

**fpserverIp** (STRING): IP address of the client.

# nblade.fpolpsinitfail events

## Nblade.fpolPsInitFail

### Severity

ERROR

### Description

This message occurs when initialization of any FPolicy persistent store resource fails. It includes free and active staging buffer pool, fname cache and dirty SVM cache.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Initialization of FPolicy resource the "%s" is unsuccessful.

### Parameters

**resource** (STRING): Uninitialized resource.

## nblade.fpolpsvolmoving events

### Nblade.fpolPsVolMoving

### Severity

ERROR

### Description

This message occurs when the FPolicy persistent store volume is being moved. This causes disruption in FPolicy where persistent store is enabled.

### Corrective Action

Disable the FPolicy with a persistent store until volume move is completed.

### Syslog Message

FPolicy persistent store volume "%s" belonging to SVM "%s" is being moved.

### Parameters

**volName** (STRING): Persistent store volume.

**vserverName** (STRING): Name of the SVM.

## nblade.fpolpsvolnoaccess events

### Nblade.fpolPsVolNoAccess

### Severity

ERROR

### Description

This message occurs when events cannot be written into persistent store volume due to insufficient permissions. This causes disruption in FPolicy where persistent store is enabled.

### Corrective Action

Check the permissions of the persistent store volume and provide full access.

### Syslog Message

Events cannot be written into the persistent store volume "%s" in SVM "%s" due to insufficient permissions.

### Parameters

**volName** (STRING): Persistent store volume

**vserverName** (STRING): Name of the SVM.

## nblade.fpolpsvolnospace events

### Nblade.fpolPsVolNoSpace

#### Severity

ERROR

#### Description

This message occurs when persistent store volume is full. This causes disruption in FPolicy where persistent store is enabled.

#### Corrective Action

Increase the persistent store volume size.

#### Syslog Message

FPolicy persistent store volume "%s" belonging to SVM "%s" has no space for new notifications. This leads to drop of notifications.

#### Parameters

**volName** (STRING): Persistent store volume.

**vserverName** (STRING): Name of the SVM.

## nblade.fpolpsvolnotfound events

### Nblade.fpolPsVolNotFound

#### Severity

ERROR

#### Description

This message occurs when volume corresponding to persistent store is not found. This causes disruption in FPolicy where persistent store is enabled.

#### Corrective Action

Check the volume corresponding to the persistent store. Either recreate the volume with same name or delete the persistent store.

#### Syslog Message

FPolicy persistent store volume "%s" belonging to SVM "%s" was not found.



## Parameters

**volName** (STRING): Persistent store volume.

**vserverName** (STRING): Name of the SVM.

# nblade.fpolpsvolofflined events

## Nblade.fpolPsVolOffline

### Severity

ERROR

### Description

This message occurs when persistent store volume is offline. This causes disruption in FPolicy where persistent store is enabled.

### Corrective Action

Check the volume corresponding to the persistent store. Either online the volume for use or delete the corresponding persistent store.

### Syslog Message

FPolicy persistent store volume "%s" belonging to SVM "%s" is offline.

## Parameters

**volName** (STRING): Persistent store volume.

**vserverName** (STRING): Name of the SVM.

# nblade.gracebegin events

## Nblade.graceBegin

### Severity

NOTICE

### Description

This message occurs when the NFS server enters the grace state.

### Corrective Action

(None).

### Syslog Message

NFS server grace state has begun for Vserver "%s", LIF ID "%d", LIF IP address "%s".

## Parameters

**vserverName** (STRING): Vserver Name.

**lifId** (INT): LIF ID

**lifIpAddress** (STRING): LIF IP address for the NFS server.

# nblade.graceend events

## Nblade.graceEnd

### Severity

NOTICE

### Description

This message occurs when the NFS server exits the grace state.

### Corrective Action

(None).

### Syslog Message

NFS server grace state has ended for Vserver "%s", LIF ID "%d", LIF IP address "%s".

### Parameters

**vserverName** (STRING): Vserver Name.

**lifid** (INT): LIF ID

**lifIpAddress** (STRING): LIF IP address for the NFS server.

# nblade.httpmaxconntimeouts events

## Nblade.httpMaxConnTimeOuts

### Severity

ERROR

### Description

This message occurs when the HTTP subsystem, which monitors connection timeouts due to an unresponsive client or data not received from the client within a specified time, especially a PUT request or an HTTP Request containing Request Body. If the number of timeouts allowed in a given period of time is being exceeded, it could be due to network failures or DoS(Denial of Service) attacks. This might result in HTTP client connection failures, because most connections are using resources and waiting for data.

### Corrective Action

Inspect the clients that are experiencing a long delay in packet trace tools. Inspecting the client network might be helpful as well.

### Syslog Message

The HTTP subsystem on this node on vserver %u has exceeded its connection timeout limit of %llu times, currently %llu connection timeout failures observed since boot time. This might result in unexpected HTTP connection failures.

### Parameters

**vsId** (INT): Id of the vserver, on which this event occurred.

**currentNumConnTimeouts** (LONGINT): Total number of connection timeouts that occurred in the HTTP subsystem.

**totalNumConnTimeouts** (LONGINT): Maximum number of connection timeouts allowed in the HTTP subsystem.

# nblade.httpmemexceeded events

## Nblade.httpMemExceeded

### Severity

ERROR

### Description

This message occurs when the HTTP subsystem, which operates all the HTTP servers exposed through a given node, has requested more memory than it is allowed in this node. The subsystem is restricted to a percentage of total memory in the node, so the node itself need not be low on memory for this to occur. HTTP memory is not allocated in advance and so HTTP memory requests compete with other system processes. This message might not appear if a memory request from HTTP cannot be honored because the node is already low on memory. This might result in HTTP client connection failures.

### Corrective Action

Compare the number of memory allocation failures with the previous occurrence of this event to understand the load on the system. Further analyze the memory allocated to the HTTP subsystem on this node by viewing the nblade\_http object instance for this node through the "statistics show" interface. Based on that data, consider distributing HTTP load by moving data LIFs, adding data LIFs, and redirecting workloads, or by shutting down non-essential workloads to alleviate the memory pressure.

### Syslog Message

The HTTP subsystem on this node has exceeded its allotment of %llu bytes of node memory with currently %llu memory allocation failures since boot time. This might result in unexpected HTTP connection failures.

### Parameters

**availableMemoryInBytes** (LONGINT): Amount of memory available to the HTTP subsystem.

**memAllocFailures** (LONGINT): Number of memory allocation failures in the HTTP subsystem since boot time.

# nblade.invtcprecordmarker events

## Nblade.invTcpRecordMarker

### Severity

ERROR

### Description

This message occurs when a TCP connection is reset due to an invalid TCP record marker in an Open Network Computing/Remote Procedure Call request.

### Corrective Action

1. Use the IP address from the log to identify the client machine that sent the invalid TCP record marker.
2. Use the TCP port number to determine which service was being requested by the client.
3. Contact the OS vendor for the client machine for a software patch.

### Syslog Message

Invalid TCP record marker received from remote IP address %s on local port %d. Vserver ID associated with this operation is %d. Invalid reason is %s, invalid record marker is %d.

## Parameters

**remoteIpAddress** (STRING): IP address of the client machine.

**localPort** (INT): Local TCP port number that received the data.

**VirtualServer** (INT): Identifier for the Vserver associated with this operation.

**reason** (STRING): Reason the record marker was invalid: - "TCP\_FRAGMENT" indicates that RPC TCP fragments that are not supported at this time. - "TCP\_DATA\_LENGTH\_OVERFLOW" indicates an oversized record marker; that is, the record marker value was greater than the maximum supported size, currently 1MB. - "TCP\_DATA\_LENGTH\_UNDERFLOW" indicates an undersized record marker; that is, the record marker value was below the minimum supported size, currently 40 bytes.

**recordMarker** (INT): Actual 32-bit record marker that was received and marked as invalid.

# nblade.junctionrootlookup2 events

## Nblade.JunctionRootLookup2

### Severity

ERROR

### Description

This message occurs when the system cannot obtain the root file handle of the volume mounted on a junction path. The cause of the problem might be that the target node is down or the target volume or aggregate is offline. The typical client operation that triggers a junction root lookup is a directory list operation where one of the directory entries is a junction. As long as the root of the volume at the junction cannot be looked up, the NFS client that is listing the directory appears to hang.

### Corrective Action

Check whether the volume or aggregate or node that matches the given MSID is offline. If it is, bring it online if possible.

### Syslog Message

Junction root lookup failed for a volume (Name: %s, MSID: %llu) in Vserver (Name: %s, ID: %llu) with reason: %s.

### Parameters

**volName** (STRING): Name of the volume mounted at the junction being looked up.

**msid** (LONGINT): Master Data Set ID (MSID) of the volume mounted at the junction being looked up.

**vserverName** (STRING): Vserver name for this MSID.

**vserverId** (LONGINT): Vserver for this MSID.

**reason** (STRING): Additional information about why the operation failed.

# nblade.longrunningrpcop events

## Nblade.longRunningRpcOp

### Severity

ERROR

### Description

This message occurs when a running remote procedure call (RPC) operation takes longer than a defined length of time. The message indicates that the RPC operation has not been completed. If the condition persists, any operation to that RPC server is likely to fail.

## Corrective Action

Check the progress of the operation by using the (privilege: advanced) "statistic show -component nblade\_rpc\_server -instance [RPC\_server\_name] -raw true -counter curr\_head\*" command. If the curr\_head\_proc\_time\_secs counter keeps incrementing, the operation has not been completed. This condition might cause network configuration and status commands to fail consistently. If the operation does not get completed, contact technical support for assistance.

## Syslog Message

Detected a long-running RPC operation (procedure number %u) for RPC server %s. The server has not responded in %lu seconds, which is over the threshold of %lu seconds.

## Parameters

**procedure** (INT): Procedure number of the RPC operation.

**RpcServer** (STRING): Name of the RPC server.

**timeSecs** (LONGINT): Time in seconds since the current longest operation was allocated.

**timeThresholdSecs** (LONGINT): Time threshold that an operation should finish by.

# nblade.nbnstoomanylifs events

## Nblade.NbnsTooManyLifs

### Severity

ERROR

### Description

This message occurs when the NetBIOS Name Service receives more than 25 IP addresses for a name registration request. This is typically caused because of an upgrade from a previous release that had more than 25 IP addresses configured for a Vserver and WINS was configured. In such a scenario, clients might not be able to access data because some of the IP addresses might not have registered with WINS.

### Corrective Action

If there are more than 25 IP addresses configured for a Vserver, modify the "nbns-enabled-lifs" CIFS option to no more than a maximum of 25 LIFs by using the 'vserver cifs options modify -vserver <vserver> -nbns -enabled-lifs' command in advanced privilege mode when the effective cluster version is Data ONTAP 9.0 or later. In a mixed-mode cluster with a node running Data ONTAP 8.3, do not configure more than 25 LIFs per Vserver because of the following WINS server limitation: [https://www.microsoft.com/resources/documentation/windowsnt/4/server/reskit/en-us/net/sur\\_wins.msp](https://www.microsoft.com/resources/documentation/windowsnt/4/server/reskit/en-us/net/sur_wins.msp)

### Syslog Message

The NetBIOS Name Service received more than 25 IP addresses for name registration for Vserver %u.

### Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

# nblade.newclientidmismatch events

## Nblade.NewClientIdMismatch

### Severity

ALERT

## Description

This message occurs when a namestring collision occurs between NFS clients.

## Corrective Action

Ensure client has unique hostname and NFSv4 mount clientaddr to avoid this.

## Syslog Message

NFSv4 name string "%s" collision between clients %s and %s.

## Parameters

**clientNameStr** (STRING): Client identifier namestring during SETCLIENTID operation.

**existingStr** (STRING): IP address of the existing client.

**newStr** (STRING): IP address of the new client.

# nblade.nfs4illegaldirentname events

## Nblade.Nfs4IllegalDirentName

### Severity

ERROR

### Description

This message occurs when a user attempts to list a directory over NFSv4 that contains an entry name with non UTF-8 characters. The problematic directory entry name was most likely created over NFSv2 or NFSv3 and could not be translated to Unicode using the translation table determined by the virtual server's language setting. A possible reason for this could be a client or server language misconfiguration issue.

### Corrective Action

Use the "vol show -fields msid" command to find the volume name corresponding to the volume MSID. Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Access the parent directory from an NFSv3 client, and then rename the entry using Unicode characters.

### Syslog Message

Directory entry with file ID '%u' on the volume with MSID '%llu' cannot be listed. Directory entry name is illegal for NFSv4.

### Parameters

**fileid** (INT): File ID of the directory entry.

**msid** (LONGINT): MSID of the volume, where the directory entry resides.

# nblade.nfs4opnotsupported events

## Nblade.Nfs4OpNotSupported

### Severity

NOTICE

### Description

This message occurs when a NFSv4 client presents an unsupported operation to the server.

## Corrective Action

(None).

## Syslog Message

NFSv4 operation '%d' is not supported.

## Parameters

**opcode** (INT): ID of unsupported operation.

# nblade.nfs4sequenceinvalid events

## nblade.nfs4SequenceInvalid

### Severity

NOTICE

### Description

This message occurs when an NFS server detects an unexpected value for an incoming sequence number from a client.

### Corrective Action

(None).

### Syslog Message

NFS client (IP: %s) sent sequence# %d, but server expected sequence# %d. Server error: %s.

### Parameters

**clientIpAddress** (STRING): IP address of the NFS client.  
**InSeqid** (INT): Incoming sequence number from the client.  
**NextSeqid** (INT): Expected sequence number from the server.  
**NfsError** (STRING): NFS protocol error.

# nblade.nfs4trunkmaxlimit events

## Nblade.nfs4TrunkMaxLimit

### Severity

ERROR

### Description

This message occurs when a session trunking group created by the NFSv4.1 client exceeds the maximum number of supported trunk connections. NFS traffic does not respond on the new trunk connection over the maximum limit.

### Corrective Action

Use only up to the maximum number of supported LIFs for session trunked NFS4.1 clients. Modify the NFS client configurations to limit the number of trunked sessions to the maximum supported to avoid the issue.

## Syslog Message

In-flight requests from client IP [%s]:%d to destination LIF "%s" (Vserver: %s) is trying to use more than max supported:"%d" trunking connections. This results in NFS41 request getting not responded by NFS server due to limit reached.

## Parameters

**clientIpAddress** (STRING): IP address of the client.

**clientPort** (INT): Network port of the client.

**lifIpAddress** (STRING): IP address of the logical interface.

**vserverName** (STRING): Name of the Vserver associated with this operation.

**maxTrunkConnectionLimit** (INT): Maximum number of allowed trunked connection by NFS41 server.

# nblade.nfs4unoptimizedio events

## Nblade.nfs4UnOptimizedIO

### Severity

ALERT

### Description

This message occurs when a session trunking group that is created by the NFSv41 client is split across multiple nodes. This can occur when one or more trunked data LIFs used in the trunking group migrate to the partner node. A split session trunking group can severely degrade performance of NFS client operations.

### Corrective Action

Address any network or hardware issues that might have caused the LIFs to migrate. Ensure trunked data LIFs are on the same node. Use the "network interface show" and "network interface migrate" commands to revert the migrated LIF to the original node where the trunking relationship was created for that LIF.

## Syslog Message

In-flight requests from client IP [%s]:%d to destination LIF "%s" (Vserver: %s) was using an unoptimized path for processing the request because of a split session trunking group. The NFSv4 client might experience degraded performance.

## Parameters

**clientIpAddress** (STRING): IP address of the client.

**clientPort** (INT): Network port of the client.

**lifIpAddress** (STRING): IP address of the logical interface.

**vserverName** (STRING): Name of the Vserver associated with this operation.

# nblade.nfsconnresetandclose events

## Nblade.nfsConnResetAndClose

### Severity

ERROR

### Description

This message occurs when the server aborts the connection with the client by sending one or more RST segments, immediately resulting in the connection state being discarded.



## Corrective Action

If this message occurs more than three times for a particular client in 30-60 minutes, make note of the IP address of the affected client. Verify that the network between the client and the storage system is functioning normally. After verifying the health of the network, verify that there are no issues related to networking on the client side. Check IP connectivity over the cluster interfaces between the two nodes. If the problem persists after addressing issues on the network or the client, contact NetApp technical support.

## Syslog Message

Shutting down connection with the client. Vserver ID is %d; network data protocol is %s, Rpc Xid 0x%x; client IP address:port is %s:%d. local IP address is %s; reason is %s.

## Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**serviceProtocol** (STRING): Network data protocol used in the connection.

**rpcXid** (INT): Remote procedure call XID.

**remoteAddr** (STRING): IP address of the client machine.

**remotePort** (INT): TCP or UDP port number used by the remote client to send the data.

**localAddr** (STRING): IP address of the local interface serving the protocol operation.

**reason** (STRING): Reason for the disconnection.

# nblade.nfscredcacheflushed events

## Nblade.nfsCredCacheFlushed

### Severity

INFORMATIONAL

### Description

This message occurs when the administrator modifies the "extended-groups-limit" option or "auth-sys-extended-groups" option using the "vserver nfs modify" command. This results in the flushing of the entire credential cache, thereby making the subsequent operations slower for a short while until the credential cache is repopulated.

### Corrective Action

(None).

### Syslog Message

When the administrator modifies the "extended-groups-limit" option or "auth-sys-extended-groups" option using the "vserver nfs modify" command, the entire credential cache is flushed that holds credentials on connections that use mixed-mode security style volumes or RPCSEC\_GSS authentication or extended groups over AUTH\_SYS. This makes subsequent operations on such connections slower for a short while, until the credential cache is repopulated. The value of "auth-sys-extended-groups" option is %d (1:enabled, 0:disabled). The value of "extended-groups-limit" option is %d.

### Parameters

**isExtGroupsEnabled** (INT): Flag to indicate if "auth-sys-extended-groups" option is enabled or not.

**extendedGroupsLimit** (INT): This indicates the number of auxiliary groups supported over RPC security flavors AUTH\_SYS or RPCSEC\_GSS. The range is 32 to 1024.

# nblade.nfsmountrootonly events

## Nblade.NfsMountRootOnly

### Severity

NOTICE

### Description

This message occurs when a non-root user performs MOUNT or NFS operation when the respective options '-mount-rootonly' or 'nfs-rootonly' are enabled.

### Corrective Action

(None).

### Syslog Message

%s operation by non-root user failed. Vserver ID is %d; ProgramNumber is %d; client IP address:port is %s:%d; local IP address is %s; Procedure number is %d.

### Parameters

**serviceProtocol** (STRING): Network data protocol used in the connection.

**vserverId** (INT): Identifier for the vsver associated with this operation.

**programNumber** (INT): RPC Program Number.

**remoteAddr** (STRING): IP address of the client machine.

**remotePort** (INT): Remote port number that sent the data.

**localAddr** (STRING): IP address of the local interface serving the protocol operation.

**procedureNum** (INT): RPC Procedure Number for the requested operation.

# nblade.nfspathresmaxlinks events

## Nblade.nfsPathResMaxLinks

### Severity

ERROR

### Description

This message occurs while resolving an input path, when the server has reached the limit of maximum links it can follow and the corresponding RPC fails.

### Corrective Action

Use the client IP address included in this message to identify the client sending the request. An operation such as mount or quota. failing on this client might be due to this error. Reduce the number of symlinks to be followed for the failing operation.

### Syslog Message

NFS maximum symbolic link limit reached when resolving path. Vserver ID: %d, service protocol:version is %s:%d, client IP address:port is %s:%d, local IP address: %s, RPC procedure number: %d.

### Parameters

**vserverId** (INT): Identifier for the Vserver associated with this operation.

**serviceProtocol** (STRING): Network data protocol used in the connection.

**programVersion** (INT): RPC program version.

**remoteAddr** (STRING): IP address of the client machine.  
**remotePort** (INT): Remote port number that sent the data.  
**localAddr** (STRING): IP address of the local interface serving the protocol operation.  
**procedureNum** (INT): RPC procedure number for the requested operation.

## nblade.nfsrcchecksummismatch events

### Nblade.nfsRCChecksumMismatch

#### Severity

ALERT

#### Description

This message occurs when the server detects a retransmitted request from the client with a modified payload. Modified requests from the client could cause corruption and pose a security threat to the server. Please use the information in the event to identify and monitor the bad client connection.

#### Corrective Action

1. Use the IP address from the log to identify the client machine and connection that retransmitted the invalid modified request. 2. Monitor the client connection and if invalid modified requests are persistent, take steps to ensure the security of the connection has not been compromised. 3. If there are no security concerns with the client connection, then the application driving traffic on that connection and the client OS NFS implementation should be examined to identify potential software defects.

#### Syslog Message

Replay cache checksum mismatch detected in retransmitted NFS request from client. Vserver ID is %d; service protocol:version is %s:%d, Rpc Xid 0x%x; client IP address:port is %s:%d. local IP address is %s; NFS procedure number is %d.

#### Parameters

**vserverId** (INT): Identifier for the vservers associated with this operation.  
**serviceProtocol** (STRING): Network data protocol used in the connection.  
**programVersion** (INT): RPC Program Version.  
**rpcXid** (INT): Remote procedure call XID.  
**remoteAddr** (STRING): IP address of the client machine.  
**remotePort** (INT): Remote port number that sent the data.  
**localAddr** (STRING): IP address of the local interface serving the protocol operation.  
**procedureNum** (INT): RPC Procedure Number for the requested operation.

## nblade.nfsv3writetoolarge events

### Nblade.Nfsv3WriteTooLarge

#### Severity

ERROR

#### Description

This message occurs when a client attempts to write over NFSv3 with an amount of data that is greater than the maximum allowed TCP or UDP transfer size. This is an invalid request from the client and the data is not written.

## Corrective Action

Try the operation again with the number of bytes of data that is less than or equal to the maximum allowed TCP or UDP transfer size configured in the server. You can set the write size option in the 'mount' command to the TCP or UDP transfer size.

## Syslog Message

%d bytes of data to be written is greater than the maximum allowed TCP or UDP transfer size, which is %d. Data is not written. Associated object type is %s with UUID: %s.

## Parameters

**opaqueLen** (INT): Number of bytes of data that were to be written.  
**maxXferSize** (INT): Maximum allowed TCP/UDP transfer size configured in the server.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.nfsv4blockclient events

## Nblade.nfsV4BlockClient

### Severity

NOTICE

### Description

This message occurs when an NFSv4 client consumes the maximum allowed store pool resources and the store pool system threshold has also been reached. In response, the NFSv4 server disconnects the client, places the client in the blocked client list and does not allow it to reconnect for 24 hours.

### Corrective Action

If you need further assistance, contact NetApp technical support.

### Syslog Message

NFSv4 client "%s" has consumed the maximum store pool resources of type "%s". It is disconnected and is marked as a blocked client.

### Parameters

**clientIpAddress** (STRING): IP address of the client.  
**poolName** (STRING): NFSv4 store pool type.

# nblade.nfsv4nsdbdomainmismatch events

## Nblade.Nfsv4NsdbDomainMismatch

### Severity

ERROR

### Description

This message occurs when an NFSv4 client queries or updates the user or group attribute of a file object but the requested value does not match the '-v4-id-domain' option value of a virtual server's NFS service. In NFSv4, users and groups are represented as names of the form string@domain. The domain string must match on the client and server to ensure correct operation. The NFSv4 client request, on whose behalf this

event message is issued, is rejected with an NFS4ERR\_BADOWNER error.

### Corrective Action

Update the client or server NFSv4 domain string to ensure that they match. If a single client experiences the issue, you typically just need to modify that client. See the client vendor's documentation to change the NFSv4 domain name configuration. If multiple clients experience the issue, you typically need to update the value of the '-v4-id-domain' option in the NFS server configuration.

### Syslog Message

NFSv4 server %s received domain string %s from client %s, which does not match the '-v4-id-domain' value %s.

### Parameters

**serverIp** (STRING): IP address of the server processing the NFSv4 request.

**clientDomain** (STRING): NFSv4 domain string received from the client

**clientIp** (STRING): IP address of the client.

**serverDomain** (STRING): Domain string that the NFSv4 server expected.

## nblade.nfsv4poolexhaust events

### Nblade.nfsV4PoolExhaust

#### Severity

EMERGENCY

#### Description

This message occurs when one of the NFSv4 store pools is exhausted.

#### Corrective Action

If the NFS server is unresponsive for more than 10 minutes after this error occurs, contact NetApp technical support.

#### Syslog Message

NFS Store Pool for %s exhausted. Associated object type is %s with UUID: %s.

#### Parameters

**poolname** (STRING): NFSv4 store pool type where the exhaustion happened.

**object\_type** (STRING): Type of the resource object under notification.

**object\_uuid** (STRING): UUID of the resource object.

## nblade.nfsv4poolthreshold events

### Nblade.nfsV4PoolThreshold

#### Severity

NOTICE

#### Description

This message occurs when one of the NFSv4 store pools is nearing exhaustion (80% of max size). If the condition is allowed to continue, the store pool limits will be reached leading certain NFSv4 request failures.

### Corrective Action

Reduce NFSv4 workload like active opens, users and client connections which consume store pool objects to ensure their counts remain under pre-defined limits.

### Syslog Message

NFS Store Pool for %s is nearing exhaustion (%s of pool currently in use).

### Parameters

**poolname** (STRING): NFSv4 store pool type where the maximum limits are nearing.

**percent** (STRING): Percentage of store pool objects of type currently in use.

## nblade.nfsv4writetoolarge events

### Nblade.Nfsv4WriteTooLarge

#### Severity

ERROR

#### Description

This message occurs when a client attempts to write over NFSv4 with an amount of data that is greater than the maximum allowed TCP transfer size. This is an invalid request from the client and the data is not written.

#### Corrective Action

Try the operation again with the number of bytes of data that is less than or equal to the maximum allowed TCP transfer size configured in the server. You can set the write size option in the 'mount' command to the TCP transfer size.

#### Syslog Message

%d bytes of data to be written is greater than the maximum allowed TCP transfer size, which is %d. Data is not written. Associated object type is %s with UUID: %s.

#### Parameters

**opaqueLen** (INT): Number of bytes of data that were to be written.

**maxtcpxferSize** (INT): Maximum allowed TCP transfer size configured in the server.

**object\_type** (STRING): Type of the resource object under notification.

**object\_uuid** (STRING): UUID of the resource object.

## nblade.nocsmsession events

### Nblade.noCsmSession

#### Severity

NOTICE

#### Description

This message occurs when the local node could not acquire a Connection Session Manager (CSM) session. CSM sessions are used to communicate file operations between nodes in the cluster. This might be a transient condition. If this condition persists, volumes on the destination node remain unavailable from this node. The destination node can be either local or remote.

### Corrective Action

Check IP connectivity over the cluster interfaces between the two nodes. If further assistance is needed, contact NetApp technical support.

### Syslog Message

(None).

### Parameters

**bladeId** (STRING): UUID of the node that could not be reached.

## nblade.nosmbvernegotiated events

### Nblade.noSmbVerNegotiated

#### Severity

ERROR

#### Description

This message occurs when Data ONTAP® does not support any of the SMB protocol versions sent by the client for negotiation.

#### Corrective Action

Ensure that the client supports at least one of the enabled SMB versions. You can determine whether an SMB version is enabled by running the "vserver cifs options show" command in advanced mode.

#### Syslog Message

For Vserver "%s", Data ONTAP® could not negotiate with the client "%s" because it does not support any of the SMB protocol versions that the client sent in the list "%s".

#### Parameters

**vserverName** (STRING): Vserver associated with this operation.

**clientIp** (STRING): IP address of the client attempting to negotiate.

**clientVersionList** (STRING): List of the SMB protocol versions sent by the client.

## nblade.notactivevscanserver events

### Nblade.NotActiveVscanServer

#### Severity

ERROR

#### Description

This message occurs when a connection from a remote IP is not found in the Vscan scanner pool belonging to the storage virtual machine (SVM). Session setup request fails.

#### Corrective Action

Investigate if the associated IPs need to be added to the active scanner pool.

## Syslog Message

CIFS connection on SVM %s (ID %u) from IP address %s fails to bind to an established Vscan related session because the IP address is not in the active scanner pool.

## Parameters

**vserverName** (STRING): Name of the SVM associated with this operation.  
**vsId** (INT): ID of the SVM on which this event occurred.  
**remotelpAddress** (STRING): IP address of the client machine.

# nblade.pcp events

## Nblade.pcp.linkThrashing

### Severity

INFORMATIONAL

### Description

This message occurs when a network port undergoes multiple link state transitions within a short period of time. The event might be triggered in response to a rapid sequence of switch events or there might be a bad link.

### Corrective Action

If this problem persists, check the physical connections between the affected port and its network switch. You can identify the faulty port by reviewing the event log for link state transitions reported by the NIC. One or more ports should have reported multiple linkUp and linkDown events prior to the linkThrashing message. ::> event log show -node <node> -messagename netif.link\* If checking physical connections does not identify the problem, consider replacing the corresponding NIC.

## Syslog Message

Link thrashing detected on port %u.

## Parameters

**port** (INT): Port ID.

# nblade.qosvioldetectregfail events

## Nblade.QoSViolDetectRegFail

### Severity

ERROR

### Description

This message occurs when the Quality-of-Service (QoS) subsystem fails to register for updates during initialization due to an internal error (e.g. there is not enough system memory available). As a result, workloads might not meet their minimum throughput or their service-level objectives (SLOs).

### Corrective Action

The registration will continue to be retried. If this message continues to occur, perform a planned takeover/giveback sequence to reset the reporting node nondisruptively by using the "takeover -ofnode (reporting\_node)" and "giveback -ofnode (reporting\_node)" commands. If the problem persists, contact Contact NetApp technical support.



## Syslog Message

Failed to register for violation detection updates. QoS min-throughput might operate in a limited capacity.

## Parameters

**errorCode** (INT): Error code that was returned during registration.

# nblade.rcbinid events

## nblade.RCBinId.timeout

### Severity

NOTICE

### Description

This message occurs when an NFS replay cache bin ID entry expires.

### Corrective Action

(None).

## Syslog Message

Found replay cache expired entry for NFS client IP address: %s.

## Parameters

**ipaddr** (STRING): IP address of the NFS client.

# nblade.recoverybegin events

## Nblade.recoveryBegin

### Severity

NOTICE

### Description

This message occurs when the NFS server enters the lock recovery state.

### Corrective Action

(None).

## Syslog Message

NFS server lock recovery has begun for Vserver "%s", LIFID "%d", LIF IP address "%s".

## Parameters

**vserverName** (STRING): Vserver Name.

**lifid** (INT): LIF ID

**lifIpAddress** (STRING): LIF IP address for the NFS server.

# nblade.s3bucketpathnotfound events

## Nblade.s3BucketPathNotFound

### Severity

ERROR

### Description

This message occurs when the path of the NAS directory corresponding to an S3 NAS Bucket was not found. This would cause all accesses to that S3 bucket to fail.

### Corrective Action

Check the NAS path corresponding to the S3 bucket. Either recreate the NAS path or delete the S3 bucket.

### Syslog Message

NAS path "%s" corresponding to S3 bucket "%s" was not found on Vserver "%s".

### Parameters

**nasPath** (STRING): NAS directory path of the S3 bucket that was not found.

**bucketName** (STRING): Name of the S3 bucket associated with this operation.

**vserverName** (STRING): Vserver associated with this operation.

## nblade.scannerconnected events

### Nblade.scannerConnected

### Severity

NOTICE

### Description

This message occurs when a Vserver connects successfully to a Vscan server.

### Corrective Action

(None).

### Syslog Message

Vserver "%s" connected to Vscan server (IP: %s).

### Parameters

**vserverName** (STRING): Name of the connecting Vserver.

**scannerIp** (STRING): IP address of the Vscan server.

## nblade.scannerdisconnected events

### Nblade.scannerDisconnected

### Severity

NOTICE

### Description

This message occurs when any Vserver disconnects from a Vscan server.

### Corrective Action

(None).

### Syslog Message

Vserver "%s" disconnected from Vscan server (IP: %s).

### Parameters

**vserverName** (STRING): Name of the disconnected Vserver.

**scannerIp** (STRING): IP address of the Vscan server.

## nblade.sectracequeueoverload events

### Nblade.sectraceQueOverload

#### Severity

ALERT

#### Description

This message occurs when there are too many Sectrace work items. This might be the result of a very generic filter. For example, tracing of succesful NFS accesses can quickly generate a large number of Sectrace items, since all ops (each read or write) might generate a security trace.

#### Corrective Action

Check the Sectrace filter and make it more specific.

#### Syslog Message

Too many events [%d] present in SectraceGlobal task queue.

#### Parameters

**queueSize** (INT): Number of queued events waiting to be processed.

## nblade.sharemaxconnlimit events

### Nblade.shareMaxConnLimit

#### Severity

ERROR

#### Description

This message occurs when a CIFS share exceeds its maximum configured connection limit.

#### Corrective Action

Increase the "max-connections-per-share" limit using "vserver cifs share modify" command.

#### Syslog Message

CIFS share: %s of Vserver: %s hit its maximum connection limit.

#### Parameters

**shareName** (STRING): Name of the CIFS share that the user attempted to connect.

**vserverName** (STRING): Name for the Vserver associated with this operation.

# nblade.showmountdisabled events

## nblade.showmountdisabled

### Severity

ERROR

### Description

This message occurs when the "showmount" command is triggered by client, but the command is disabled on the NFS server. The "showmount" command enables NFS clients to see the Vserver's NFS exports list.

### Corrective Action

Use the "nfs modify -showmount enabled" command to enable the feature. Use the "nfs show -fields showmount" command to verify that, if the showmount feature was enabled successfully on the NFS server.

### Syslog Message

Showmount is disabled on Vserver %d.

### Parameters

**vserverid** (INT): Identifier for the Vserver associated with this operation.

# nblade.smb1onlyclientfound events

## Nblade.smb1OnlyClientFound

### Severity

NOTICE

### Description

This message occurs when ONTAP® detects a client that supports only the SMB1 protocol. This message can help in identifying such clients and preparing for their upgrade to SMB2 or later before ONTAP deprecates the SMB1 version of the protocol.

### Corrective Action

(None).

### Syslog Message

For Vserver "%s", Data ONTAP® detected a client "%s" that supports only the SMB1 version of the protocol.

### Parameters

**vserverName** (STRING): Vserver associated with this operation.

**clientIp** (STRING): IP address of the client.

# nblade.smbsignaturemismatch events

## Nblade.smbSignatureMismatch

**Severity**

ERROR

**Description**

This message occurs when the SMB signature sent by client does not match the SMB signature generated by the CIFS server.

**Corrective Action**

Ensure that the SMB client is authenticated in order to access the files. Alternatively, disable SMB signing.

**Syslog Message**

Client (protocol dialect: %s, IP: %s) sent an SMB signature that does not match the signature generated by the CIFS server: %s (Vserver: %s).

**Parameters**

**SmbDialect** (STRING): SMB protocol dialect negotiated by the client.

**IpAddress** (STRING): IP address of the client.

**CifsServerName** (STRING): Cifs server name of the client.

**VserverName** (STRING): Vserver name.

## nblade.testevent events

**Nblade.TestEvent****Severity**

INFORMATIONAL

**Description**

This is a test EMS event generated by the Network Blade for functional verification only. No action is required.

**Corrective Action**

No corrective action is required. If this message is appearing without being intentionally triggered by the customer or by support activity, please contact support.

**Syslog Message**

The Network Blade kernel module has generated a test EMS event.

**Parameters**

(None).

## nblade.vifdel events

**Nblade.vifdel.ipRemoveFailed****Severity**

ERROR

## Description

This message occurs when there is a failure removing the IP address of a LIF that is being deleted. While this condition persists, most networking configuration or status operations can not be performed.

## Corrective Action

A system reboot may be required if the operation did not complete successfully after it was retried. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. After the reboot is complete, issue a "giveback" command to return services to this partner.

## Syslog Message

Error deleting IP address from LIF %d, error %d.

## Parameters

**lif** (INT): LIF Id.

**error** (INT): Operation error code.

## Nblade.vifdel.longRunning

### Severity

ERROR

### Description

This message occurs when there is an excessive delay in removing connections before LIF deletion can be completed. While this condition persists, most networking configuration or status operations can not be performed.

### Corrective Action

A system reboot may be required if this message has occurred multiple times for the same LIF and the number of connections remaining has not decreased across these messages. In addition, network configuration and status commands are still not working. For HA configurations, operational disruption can be minimized by initiating a partner takeover followed by a reboot of this partner. Prior to reboot or takeover, a full autosupport must be generated for the affected node. This is done via the "autosupport invoke -type full" command. If possible, a sync core of the affected node should be provided to help find the root cause. A sync core can be generated via the "reboot -dump true" command for the affected node. After the reboot is complete, issue a "giveback" command to return services to this partner for HA configurations.

### Syslog Message

Excessive delay while deleting connections on LIF %d, %d connections remaining.

### Parameters

**lif** (INT): LIF Id.

**remainingConns** (INT): Number of connections still remaining on this LIF.

## nblade.vldb events

### Nblade.vldb.LeakedBladeList

### Severity

NOTICE

## Description

This message occurs when a protocol requests a list of blades but one cannot be generated due to one or more references to a previous request or requests still outstanding.

## Corrective Action

If this condition is not cleared by a subsequent free message within 30 minutes, a reboot of the node might be required. Rebooting the node with the "system node reboot -dump true" command saves a core dump that allows diagnosis of this issue by NetApp technical support. For more information or assistance, contact NetApp technical support.

## Syslog Message

"%s" operation issued by "%s" on cached volume location database (VLDB) blade list with references pending beyond expiry time.

## Parameters

**operation** (STRING): Type of operation the caller is attempting.  
**caller** (STRING): Caller that is trying to perform the operation.

## Nblade.vldb.Timeout

### Severity

ERROR

### Description

This message occurs when a request from the network module to the volume location database (VLDB) times out. The VLDB is used to identify the node that currently controls a volume. This might be a transient condition. If the condition persists, some volumes might be inaccessible.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Request from the network module to the volume location database (VLDB) timed out.

### Parameters

(None).

## nblade.vscanbadipprivaccess events

### Nblade.vscanBadIPPrivAccess

#### Severity

ERROR

#### Description

This message occurs during a client attempt to connect to the privileged ONTAP\_ADMIN\$ share, when the IP address of that client is not found in the list of allowed IP addresses.

#### Corrective Action

Ensure that the mentioned IP address is configured in one of the active vscan scanner pools. Use the "vserver vscan scanner pool show-active" command to view the currently active scanner pool configuration.

## Syslog Message

For Vserver "%s", the attempt to connect to the privileged ONTAP\_ADMIN\$ share by the client "%s" is rejected because its IP or hostname is not configured in any of the Vserver active scanner pools.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the client.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanbadprotomagicnum events

## Nblade.vscanBadProtoMagicNum

### Severity

ERROR

### Description

This message occurs when an incorrectly formatted message is received from the AV Connector. This might result in delay or failure in AV scanning.

### Corrective Action

Ensure that the compatible version of AV Connector is running on the AV server, and that no other user or software is attempting to connect to the "\\PIPE\vscan" resource on the Vserver.

## Syslog Message

Data ONTAP® received incorrectly formatted message for Vserver "%s" from AV Connector running on the AV server "%s".

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the AV server.

# nblade.vscanbaduserprivaccess events

## Nblade.vscanBadUserPrivAccess

### Severity

ERROR

### Description

This message occurs during a client attempt to connect to the privileged ONTAP\_ADMIN\$ share, when the logged-in user of that client is not found in the list of allowed users.

### Corrective Action

Ensure that the mentioned user name and IP address is configured in one of the active vscan scanner pools. Use the "vserver vscan scanner pool show-active" command to view the currently active scanner pool configuration.



## Syslog Message

For Vserver "%s", the attempt to connect to the privileged ONTAP\_ADMIN\$ share by the client "%s" is rejected because its logged-in user "%s" is not configured in any of the Vserver active scanner pools.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the client.  
**userName** (STRING): Logged-in user of the client.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanconnbackpressure events

## Nblade.vscanConnBackPressure

### Severity

ERROR

### Description

This message occurs when an AV server is too busy to accept new scan requests.

### Corrective Action

If this message occurs frequently, ensure that there are enough AV servers to handle the virus-scanning load being generated by the mentioned Vserver.

## Syslog Message

For Vserver "%s", AV server "%s" is too busy to accept new scan requests.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the AV server.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanconninactive events

## Nblade.vscanConnInactive

### Severity

NOTICE

### Description

This message occurs when ONTAP® detects a nonresponsive antivirus (AV) server and forcibly closes its vscan connection.

### Corrective Action

Ensure that the AV Connector, installed on the AV server, can connect to the Vserver and receive the scan requests.

## Syslog Message

For Vserver "%s", ONTAP® forcibly closed the vscan connection originated from the nonresponsive AV server "%s".

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the AV server.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanconninvaliduser events

## Nblade.vscanConnInvalidUser

### Severity

NOTICE

### Description

This message occurs during a client attempt to establish a vscan connection, when the logged-in user of that client is not found in the list of allowed users.

### Corrective Action

Ensure that the mentioned user name is configured in one of the active vscan scanner pools. Use the "vserver vscan scanner pool show-active" command to view the currently active scanner pool configuration.

## Syslog Message

For Vserver "%s", the vscan connection request coming from client "%s" is rejected because the logged-in user "%s" is not configured in any of the Vserver active scanner pools.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the client attempting to establish a vscan connection.  
**userName** (STRING): User name of the client attempting to establish a vscan connection.

# nblade.vscanconnreqonsmb1 events

## Nblade.vscanConnReqOnSMB1

### Severity

ERROR

### Description

This message occurs during a client attempt to establish a vscan connection over SMB1, which is not supported.

### Corrective Action

Verify that both the client and Data ONTAP® support and are configured for SMB2 or later.

## Syslog Message

For Vserver "%s", the vscan connection request coming from the client "%s" is rejected because it is not supported for SMB1.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the client.

# nblade.vscanexcessivetos events

## Nblade.vscanExcessiveTOs

### Severity

ERROR

### Description

This message occurs when ONTAP logs excessive timed-out scanning events which are due to delays in scanning response from the virus scanning (Vscan) server. Too many timed-out scanning events indicate that files are not being scanned in a timely manner.

### Corrective Action

Investigate the Vscan server for timed-out scanning events to ascertain the root cause for the delays.

## Syslog Message

Vscan timed-out scanning events exceeded %d in the last 30 minutes for Vscan server (IP: %s) in SVM "%s".

## Parameters

**ExcessiveVscanTimeOutsCount** (INT): Threshold count for for the maximum number of timed-out scanning events that can occur within a 30 minute interval.  
**scannerIp** (STRING): IP address of the Vscan server.  
**vserverName** (STRING): Name of the SVM.

# nblade.vscannodispatcher events

## Nblade.vscanNoDispatcher

### Severity

ALERT

### Description

This message occurs when the vscan-dispatcher component cannot be created. This might be due to internal errors on the system, such as nonavailability of memory. This might cause data unavailability if the scan-mandatory option is enabled.

### Corrective Action

Disable Vscan and enable it again by using the "vserver vscan disable/enable" command. If the problem persists, contact NetApp technical support.

## Syslog Message

Vscan dispatcher component for Vserver "%s" could not be created. Error: %d.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

**errorCode** (INT): Error code that describes the reason for the failure.

# nblade.vscannopolicyenabled events

## Nblade.vscanNoPolicyEnabled

### Severity

ALERT

### Description

This message occurs when a file access is not considered for virus scanning because none of the configured On-Access policies are enabled for the Vserver.

### Corrective Action

Enable one of the configured On-Access policies for the Vserver.

## Syslog Message

For Vserver "%s", the file access was not considered for virus scanning because none of the configured On-Access policies are enabled.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

# nblade.vscannoregdscanner events

## Nblade.vscanNoRegdScanner

### Severity

ERROR

### Description

This message occurs when AV Connector notifies ONTAP® that it does not have a registered scan-engine. This could happen if the scan-engine never registered or unregistered because of some reason like it is misconfigured or shutdown. This might cause data unavailability if the scan-mandatory option is enabled.

### Corrective Action

Ensure that the scan-engine software installed on the AV server is compatible with ONTAP®. Also ensure that it is running and is configured to connect to the AV Connector over local loopback.

## Syslog Message

For Vserver "%s", AV Connector running on the AV server "%s" does not have a registered scan-engine to it.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**scannerIp** (STRING): IP address of the AV server.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscannoscannerconn events

## Nblade.vscanNoScannerConn

### Severity

EMERGENCY

### Description

This message occurs when ONTAP® has no vscan connection for servicing virus scan-requests. This might cause data unavailability if the scan-mandatory option is enabled.

### Corrective Action

Ensure that the scanner pool is properly configured and the AV servers are active and connected to ONTAP.

### Syslog Message

Vserver "%s" has no virus scanner connection.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanvirusdetected events

## Nblade.vscanVirusDetected

### Severity

ERROR

### Description

This message occurs when a vscan server reports an error to the storage system. Normally this indicates that a virus has been found by the vscan server; however, other error conditions on the vscan server can result in this event. Client access to the file is denied. The vscan server might, depending on its settings and configuration, clean the file, quarantine it, or delete it.

### Corrective Action

Check the log of the vscan (antivirus) server reported in the syslog message to see if it was able to successfully quarantine or delete the infected file. If it was not able to do so, a system administrator might want to manually delete the file.

### Syslog Message

Possible virus detected. Vserver: %s, vscan server IP: %s, file path: %s, client IP: %s, SID: %s, vscan engine status: %u, vscan engine result string: %s.

## Parameters

**vserverName** (STRING): Name of the Vserver associated with this operation.  
**vscanServerIp** (STRING): IP address of the vscan server.  
**filePath** (STRING): Path of the file that was found to be infected.  
**clientIp** (STRING): IP address of the client.  
**SID** (STRING): SID of the client.  
**vscanEngineStatus** (INT): Status code returned by the vscan server.  
**vscanEngineResultString** (STRING): Result string returned by the vscan server.  
**object\_type** (STRING): Type of the resource object under notification.  
**object\_uuid** (STRING): UUID of the resource object.

# nblade.vscanworkqueueoverloaded events

## Nblade.vscanWorkQueueOverloaded

### Severity

NOTICE

### Description

This message occurs when there are too many events being generated by the Vscan subsystem in ONTAP® software, which might be the result of a misconfigured Vscan server.

### Corrective Action

Check the Vscan server configurations to make sure that there are no connectivity or configuration issues between the storage system and Vscan server. Use the "vserver vscan connection-status show-all" command to see the connection states of all configured Vscan servers.

### Syslog Message

Too many events [%d] present in Vscan work queue.

### Parameters

**workQueueDepth** (INT): Number of queued events waiting to be processed.

# ndmp events

## ndmp.ctrl events

### ndmp.ctrl.on.e0m

#### Severity

ALERT

#### Description

This message occurs when the IP address of the logical interface (LIF) that hosts the NDMP control connection is hosted on maintenance port "e0M". This port does not provide sufficient bandwidth for NDMP.

#### Corrective Action

Ensure that the LIF is moved out of port "e0M" and hosted on a high bandwidth port, or use a LIF hosted on a high bandwidth port for the control connection.

#### Syslog Message

NDMP control connection is established on "%s" which is hosted on port "%s". This is a low bandwidth port and might impact performance.

#### Parameters

**lif\_addr** (STRING): Local control connection IP address.

**lif\_port** (STRING): Port on which the control connection IP address is hosted.

## ndmp.data events

### ndmp.data.on.e0m

#### Severity

ALERT

#### Description

This message occurs when the IP address of the logical interface (LIF) that hosts the NDMP data connection is hosted on maintenance port "e0M". This port does not provide sufficient bandwidth for NDMP.

#### Corrective Action

Ensure that the LIF is moved out of port "e0M" and hosted on a high bandwidth port, or create a new LIF hosted on a high bandwidth port which NDMP can use for the data connection.

#### Syslog Message

NDMP data connection is established on "%s" which is hosted on port "%s". This is a low bandwidth port and might impact performance.

#### Parameters

**lif\_addr** (STRING): Local data connection IP address.

**lif\_port** (STRING): Port on which the data connection IP address is hosted.

# ndmp.fail events

## ndmp.fail.log

### Severity

ERROR

### Description

This event is generated when there is an error in NDMP

### Corrective Action

(None).

### Syslog Message

Error message: %s

### Parameters

**Error\_message** (STRING): The error is hinted

# ndmp.max events

## ndmp.max.warning

### Severity

NOTICE

### Description

This message occurs when the maximum number of NDMP sessions has been reached.

### Corrective Action

(None).

### Syslog Message

Cannot accept new requests because the NDMP connection limit has been reached.

### Parameters

(None).

# ndmp.password events

## ndmp.password.warning

### Severity

ERROR

### Description

This message occurs on NDMP service startup, when the password for the NDMP user account is still the default password.



**Corrective Action**

Use the 'ndmpd password' command to change the default password.

**Syslog Message**

The NDMP user account has the default password set. Use the 'ndmp password' command to change the default password.

**Parameters**

(None).

# ndo events

## ndo.cr events

### ndo.cr.canceled

#### Severity

NOTICE

#### Description

This message occurs when a controller replacement operation is canceled by the operator.

#### Corrective Action

(None).

#### Syslog Message

Controller replacement operation on node "%s" is canceled by the operator.

#### Parameters

**nodeName** (STRING): Name of the node that was being replaced when the operation was canceled.

### ndo.cr.completed

#### Severity

NOTICE

#### Description

This message occurs when a controller replacement operation has been completed successfully in the HA group.

#### Corrective Action

(None).

#### Syslog Message

Controller replacement operation completed. New node "%s" (model: %s), HA partner node "%s" (model: %s)

#### Parameters

**newNodeName** (STRING): Node name of the first new controller.

**newModelName** (STRING): Model name of the first new controller.

**haNodeName** (STRING): Node name of the new HA partner controller.

**haModelName** (STRING): Model name of the new HA partner controller.

### ndo.cr.completed.firstNode

#### Severity

NOTICE

**Description**

This message occurs when a controller replacement operation has been completed on the first node of the HA group.

**Corrective Action**

(None).

**Syslog Message**

Controller replacement operation completed on the first controller of the HA group. New node: %s (model: %s).

**Parameters**

**nodeName** (STRING): Node name of the new controller.

**modelName** (STRING): Model name of the new controller.

**ndo.cr.nso.canceled****Severity**

NOTICE

**Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been canceled.

**Corrective Action**

(None).

**Syslog Message**

NSO-based controller replacement operation was canceled from node: %s. Remote cluster node: %s, remote cluster HA partner node: %s.

**Parameters**

**initiatingClusterNodeName** (STRING): Name of the node from where controller replacement operation was canceled.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

**ndo.cr.nso.completed****Severity**

NOTICE

**Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has completed.

**Corrective Action**

(None).

## Syslog Message

NSO-based controller replacement operation completed. Initiating cluster node: %s, remote cluster node: %s (model: %s), remote cluster HA partner node: %s (model: %s).

## Parameters

**initiatingClusterNodeName** (STRING): Node name of the controller that has initiated the upgrade.  
**newTargetClusterNodeName** (STRING): Node name of the new controller.  
**newTargetClusterModelName** (STRING): Model name of the new controller.  
**newTargetClusterHApartnerNodeName** (STRING): Node name of the new HA partner controller.  
**newTargetClusterHApartnerModelName** (STRING): Model name of the new HA partner controller.

## ndo.cr.nso.mccip.canceled

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been canceled.

### Corrective Action

(None).

## Syslog Message

An NSO-based controller replacement operation was canceled in a MetroCluster IP configuration from node "%s". The remote cluster node is "%s", and the remote cluster HA partner node is "%s".

## Parameters

**initiatingClusterNodeName** (STRING): Name of the node where the controller replacement operation was canceled.  
**targetClusterNodeName** (STRING): Node name of the controller that was to be replaced.  
**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that was to be replaced.

## ndo.cr.nso.mccip.completed

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has completed.

### Corrective Action

(None).

## Syslog Message

An NSO-based controller replacement operation completed in a MetroCluster IP configuration. The initiating cluster node is "%s", remote cluster node is "%s" (model "%s"), and the remote cluster HA partner node is "%s" (model "%s").

## Parameters

**initiatingClusterNodeName** (STRING): Node name of the controller that initiated the controller replacement operation.  
**newTargetClusterNodeName** (STRING): Node name of the new controller.  
**newTargetClusterModelName** (STRING): Model name of the new controller.  
**newTargetClusterHApartnerNodeName** (STRING): Node name of the new HA partner controller.  
**newTargetClusterHApartnerModelName** (STRING): Model name of the new HA partner controller.

## ndo.cr.nso.mccip.paused

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been paused.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

An NSO-based controller replacement operation was paused in a MetroCluster IP configuration from node "%s". The remote cluster node is "%s", and the remote cluster HA partner node is "%s".

## Parameters

**initiatingClusterNodeName** (STRING): Name of the node from where the controller replacement operation was paused.  
**targetClusterNodeName** (STRING): Node name of the controller that is being replaced.  
**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that is being replaced.

## ndo.cr.nso.mccip.paused.error

### Severity

ERROR

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration was paused due to an error.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Resolve the error. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

An NSO-based controller replacement operation has been paused in a MetroCluster IP configuration on node "%s" during a task "%s" due to the following error: %s.

## Parameters

**nodeName** (STRING): Name of the node where the controller replacement operation was paused due to an error.

**taskName** (STRING): Name of the task.  
**taskReport** (STRING): Detailed error report.

## **ndo.cr.nso.mccip.paused.userInput**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been paused for operator intervention.

### **Corrective Action**

Use the "system controller replace show-details" command to check the status of the operation. Collect the requested information or perform the steps as mentioned in the "system controller replace show-details" command output. Then, use the "system controller replace resume" command to resume the replacement operation.

### **Syslog Message**

An NSO-based controller replacement operation from initiating node "%s" has been paused in a MetroCluster IP configuration. The remote cluster node "%s" and remote cluster HA partner node "%s" have been paused for operator interventio while performing a "%s" task during the "%s" phase.

### **Parameters**

**initiatingClusterNodeName** (STRING): Name of the node where the controller replacement operation was initiated.

**targetClusterNodeName** (STRING): Node name of the controller that is being replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that is being replaced.

**taskName** (STRING): Name of the task where the controller replacement operation has been paused for operator intervention.

**phaseInfo** (STRING): Name of the phase where the controller replacement operation has been paused for operator intervention.

## **ndo.cr.nso.mccip.resumed**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been resumed.

### **Corrective Action**

(None).

### **Syslog Message**

An NSO-based controller replacement operation has been resumed in a MetroCluster IP configuration from node "%s". The remote cluster node is "%s", and the remote cluster HA partner node is "%s".

## Parameters

**initiatingClusterNodeName** (STRING): Name of the node from where the controller replacement operation has been resumed.

**targetClusterNodeName** (STRING): Node name of the controller that is being replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that is being replaced.

## ndo.cr.nso.mccip.started

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster IP configuration has been started.

### Corrective Action

(None).

### Syslog Message

An NSO-based controller replacement operation has been started in a MetroCluster IP configuration. The local cluster node is "%s", the remote cluster node is "%s" (model "%s"), and the remote cluster HA partner node is "%s" (model "%s").

## Parameters

**initiatingClusterNodeName** (STRING): Node name of the controller that initiated the controller replacement operation.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterModelName** (STRING): Model name of the controller that will be replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

**targetClusterHApartnerModelName** (STRING): Model name of the HA partner controller that will be replaced.

## ndo.cr.nso.paused

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been paused.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

NSO-based controller replacement operation was paused from node: %s. Remote cluster node: %s, remote cluster HA partner node: %s.

## Parameters

**initiatingClusterNodeName** (STRING): Name of the node from where the controller replacement operation was paused.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

## ndo.cr.nso.paused.error

### Severity

ERROR

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been paused due to an error.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Resolve the error. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

NSO-based controller replacement operation has been paused on node "%s" during %s task due to an error: %s.

## Parameters

**nodeName** (STRING): Name of the node where the controller replacement operation has been paused due to an error.

**taskName** (STRING): Name of the task.

**taskReport** (STRING): Detailed error report.

## ndo.cr.nso.paused.userInput

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been paused for operator intervention.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Collect the requested information or perform the steps as mentioned in the "system controller replace show-details" command. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

NSO-based controller replacement operation from initiating node "%s" has been paused. Remote cluster node %s and remote cluster HA partner node %s are paused in %s task during %s phase for operator intervention.



## Parameters

**initiatingClusterNodeName** (STRING): Name of the node where the controller replacement operation was initiated.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterHAPartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

**taskName** (STRING): NSO-based controller replacement operation from initiating node "%s" has been paused. Remote cluster node %s and remote cluster HA partner node %s are paused in %s task during %s phase for operator intervention

**phaseInfo** (STRING): Name of the phase where the controller replacement operation has been paused for operator intervention.

## ndo.cr.nso.resumed

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been resumed.

### Corrective Action

(None).

### Syslog Message

NSO-based controller replacement operation has been resumed from node %s. Remote cluster node: %s, remote cluster HA partner node: %s.

## Parameters

**initiatingClusterNodeName** (STRING): Name of the node from where the controller replacement operation has been resumed.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterHAPartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

## ndo.cr.nso.started

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation using negotiated switchover (NSO) in a MetroCluster FC configuration has been started.

### Corrective Action

(None).

### Syslog Message

NSO-based controller replacement operation has been started. Local cluster node: %s, remote cluster node: %s (model: %s), remote cluster HA partner node: %s (model: %s).

## Parameters

**initiatingClusterNodeName** (STRING): Node name of the controller that has initiated the upgrade.

**targetClusterNodeName** (STRING): Node name of the controller that will be replaced.

**targetClusterModelName** (STRING): Model name of the controller that will be replaced.

**targetClusterHApartnerNodeName** (STRING): Node name of the HA partner controller that will be replaced.

**targetClusterHApartnerModelName** (STRING): Model name of the HA partner controller that will be replaced.

## ndo.cr.paused

### Severity

NOTICE

### Description

This message occurs when a controller replacement operation has been paused by the operator.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

Controller replacement operation on node "%s" is paused by the operator.

### Parameters

**nodeName** (STRING): Name of the node where the replacement was in progress when the operation was paused.

## ndo.cr.paused.error

### Severity

ERROR

### Description

This message occurs when the controller replacement operation has been paused due to an error.

### Corrective Action

Use the "system controller replace show-details" command to check the status of the operation. Resolve the error. Then, use the "system controller replace resume" command to resume the replacement operation.

### Syslog Message

Controller replacement operation has been paused on node "%s" during %s task due to an error: %s.

### Parameters

**nodeName** (STRING): Name of the node where the controller replacement operation has been paused due to an error.

**taskName** (STRING): Name of the task.

**taskReport** (STRING): Detailed error report.

## **ndo.cr.paused.userInput**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation has been paused for operator intervention.

### **Corrective Action**

Use the "system controller replace show-details" command to check the status of the operation. Collect the requested information or perform the steps as mentioned in the "system controller replace show-details" command. Then, use the "system controller replace resume" command to resume the replacement operation.

### **Syslog Message**

Controller replacement operation has been paused on node "%s" in "%s" task during "%s" phase for operator intervention.

### **Parameters**

**nodeName** (STRING): Name of the node where the controller replacement operation has been paused for operator intervention.

**taskName** (STRING): Name of the task where the controller replacement operation has been paused for operator intervention.

**phaseInfo** (STRING): Name of the phase where the controller replacement operation has been paused for operator intervention.

## **ndo.cr.resumed**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation has been resumed by the operator.

### **Corrective Action**

(None).

### **Syslog Message**

Controller replacement operation has been resumed on node "%s" by the operator.

### **Parameters**

**nodeName** (STRING): Name of the node where the controller replacement operation has been resumed.

## **ndo.cr.started**

### **Severity**

NOTICE

### **Description**

This message occurs when a controller replacement operation has been started in the HA group by the operator.

**Corrective Action**

(None).

**Syslog Message**

Controller replacement operation has been started. Node "%s" (model: %s), HA partner node "%s" (model: %s).

**Parameters**

**nodeName** (STRING): Node name of the controller that will be replaced.

**modelName** (STRING): Model name of the controller that will be replaced.

**hanodeName** (STRING): Node name of the HA partner controller that will be replaced.

**hamodelName** (STRING): Model name of the HA partner controller that will be replaced.

# net events

## net.cg events

### net.cg.backpressureEntered

#### Severity

NOTICE

#### Description

This message occurs when the current connection group experiences backpressure. The client workload is distributed among connection groups. If a connection group has too much work waiting to be processed, incoming packets will be dropped. If one connection group experiences backpressure, it might indicate an imbalance that could affect the subset of clients using that connection group. If all connection groups are experiencing backpressure, it might indicate that the workload is nearing the capacity of the storage controller.

#### Corrective Action

Use the "statistics show -raw -object cgstat -counter is\_backpressured" command to monitor the connection group backpressure status. Use the "statistics show -raw -object cgstat -counter bkp\_exited" command to monitor the connection group backpressure rate. Use the "statistics show -raw -object cgstat -counter bkp\_ipv4\_drops" and "statistics show -raw -object cgstat -counter bkp\_ipv6\_drops" commands to monitor packet drops due to the connection group backpressure. Adjust the workload or increase the capacity to avoid over-utilization.

#### Syslog Message

Connection group "%d" experienced backpressure %d times.

#### Parameters

**connection\_group** (INT): Connection group which dropped packets due to backpressure.  
**count** (INT): Number of times the current connection group experienced backpressure.

### net.cg.ratelimited

#### Severity

NOTICE

#### Description

This message occurs when the rate of incoming traffic exceeds the threshold limit of the current connection group. The client workload is distributed among connection groups. If a connection group is receiving certain kinds of packets at a rate exceeding the allowed threshold, those incoming packets will be dropped. Packets subject to rate limits include UDP, ICMP, IP fragments, broadcast, and multicast. If a connection group is receiving packets at a rate that exceeds its threshold, the storage controller might be experiencing a packet storm.

#### Corrective Action

Use the "statistics show -raw -object cgstat -counter ratelimit\_icmp", "statistics show -raw -object cgstat -counter ratelimit\_bcast", "statistics show -raw -object cgstat -counter ratelimit\_frag", "statistics show -raw -object cgstat -counter ratelimit\_udp", and "statistics show -raw -object cgstat -counter ratelimit\_other" commands to monitor packets that are being dropped because they are being received by the connection group at a rate that exceeds the threshold limit. Use packet traces to determine if there is a packet storm or

network misconfiguration causing packets not intended for consumption by the storage controller to be sent to the connection group. If using NFS-over-UDP, adjust the workload or increase the capacity to avoid over-utilization.

### Syslog Message

"%s traffic received by connection group "%d" was received at a rate that exceeded the threshold limit.

### Parameters

**reason** (STRING): The type of packets the current connection group received at a rate greater than its threshold limit.

**connection\_group** (INT): Connection group which dropped packets because it received them at a rate greater than its threshold limit.

## net.cg.utilThresholdExceed

### Severity

NOTICE

### Description

This message occurs when the current connection group utilization exceeds the configured threshold. The client workload is distributed among connection groups. If one connection group utilization is much higher than the others, it might indicate an imbalance that could affect the subset of clients using that connection group. If all connection groups are experiencing high utilization, it might indicate that the workload is nearing the capacity of the storage controller.

### Corrective Action

Use the "statistics show -raw -object nwk\_ctx -counter utilization" command to monitor the connection group utilization. Adjust the workload or increase the capacity to avoid over-utilization.

### Syslog Message

Utilization of connection group "%d" (%d%%) exceeds the configured threshold (%d%%).

### Parameters

**connection\_group** (INT): Connection group whose utilization exceeded the configured threshold.

**current\_utilization** (INT): Current connection group utilization.

**configured\_threshold** (INT): Configured connection group utilization threshold.

## net.fib events

### net.fib.createfailed

### Severity

ERROR

### Description

This message occurs when a Vserver resource allocation fails. This leaves the system with an unusable Vserver, and additional Vservers cannot be created. This might also lead to system performance degradation.

### Corrective Action

Delete the Vserver created by using the "vserver delete" command. Use the diagnostic privilege "debug

smdb table vserver\_by\_name show" command to determine the Vserver name corresponding to the ID. You must delete the other associated entities, such as LIFs and volumes, before deleting a Vserver.

### **Syslog Message**

Resource allocation for Vserver (ID: %d) failed.

### **Parameters**

**vsid** (INT): Vserver ID.

## **net.fib.nearinglimit**

### **Severity**

NOTICE

### **Description**

This message occurs when Vserver resources have reached 90% and will soon be exhausted. When the resources are exhausted, additional Vservers cannot be created. The system will function normally with respect to other Vservers.

### **Corrective Action**

Check for configured Vservers by using the "vserver show" command, and delete unused Vservers by using the "vserver delete" command. You must delete the the other associated entities, such as LIFs and volumes, before deleting a Vserver.

### **Syslog Message**

Vserver resource allocation is nearing the limit.

### **Parameters**

(None).

## **net.ifconfig events**

### **net.ifconfig.addLimitExceed**

### **Severity**

ERROR

### **Description**

This message occurs when the limit for adding logical virtual local area network (VLAN) and interface group (ifgrp) interfaces is exceeded. The latest attempt to add an interface did not succeed.

### **Corrective Action**

Delete any existing unused VLAN or ifgrp interface and retry adding the new logical interface.

### **Syslog Message**

The limit for adding logical virtual local area network (VLAN) and interface group (ifgrp) interfaces is exceeded. The maximum limit per host is %d.

### **Parameters**

**maxlimitperhost** (INT): Maximum interfaces per host.

## net.ifconfig.ioctlError

### Severity

NOTICE

### Description

This event is generated when the ifconfig command encounters an error while attempting to perform an ioctl.

### Corrective Action

(None).

### Syslog Message

ioctl error encountered while configuring interface '%s' using ioctl '%s' (error string='%s', error code=%d).

### Parameters

**ifname** (STRING): The name of the interface

**ioctl** (STRING): The name of the ioctl

**errorString** (STRING): The error string

**error** (INT): The internal error code

## net.ifgrp events

### net.ifgrp.addLimitExceeded

#### Severity

ERROR

#### Description

This message occurs when the system fails to create a new interface group because the number of such groups already present on the node exceed either the number of physical ports on the node or the upper bound on number of interface groups supported.

#### Corrective Action

Delete any existing unused interface groups and retry adding the new interface group. Use the "network port ifgrp delete" and "network port ifgrp create" commands, respectively, to do so.

#### Syslog Message

ifgrp %s could not be created because the limit for logical interface groups on this node has been reached. The limit per node is equal to the number of physical ports present on the node with an upper bound of %d.

#### Parameters

**vifName** (STRING): Name of the interface group that could not be created.

**maxIfgrp** (INT): Upper bound on number of interface groups supported.

### net.ifgrp.lacp.key.mismatch

#### Severity

ERROR



## Description

This message occurs when ONTAP receives differing system-ID or keys information from the Link Aggregation Control Protocol (LACP) peer, leading to one or more interfaces not being active in the port interface group (ifgrp). This can indicate a configuration error at the LACP peer.

## Corrective Action

Search the knowledgebase of the NetApp technical support support web site for the "LACP port down due to mismatched peer keys" phrase.

## Syslog Message

Partner mismatch occurred on ifgrp "%s", port "%s". Expected by aggregate: system-priority "0x%02x", system-ID "%s", and key "0x%04x". Found at Port: system-priority "0x%02x", system-ID "%s", and key "0x%04x".

## Parameters

**ifgrpName** (STRING): Name of the ifgrp.  
**portName** (STRING): Name of the ifgrp member-port.  
**aggregateSysPriority** (INT): Ifgrp aggregate system priority.  
**aggregateSysMAC** (STRING): Ifgrp aggregate system MAC address.  
**aggregateKey** (INT): Ifgrp aggregate key.  
**portSysPriority** (INT): Ifgrp member-port system priority.  
**portSysMAC** (STRING): Ifgrp member-port system MAC address.  
**portKey** (INT): Ifgrp member-port key.

## net.ifgrp.lacp.link.active

### Severity

NOTICE

### Description

This message occurs when a link in an Link Aggregation Control Protocol (LACP) aggregate becomes active.

### Corrective Action

(None).

### Syslog Message

ifgrp %s, port %s has transitioned to the active state.

### Parameters

**aggregateName** (STRING): Name of the interface group.  
**portName** (STRING): Name of the member port.

## net.ifgrp.lacp.link.inactive

### Severity

ERROR

### Description

This message occurs when a link is operationally up, but becomes inactive in the Link Aggregation Control Protocol (LACP) aggregate with which it has been associated. This can occur due to a configuration, software, or link error, or a hardware failure.

### Corrective Action

Search kb.netapp.com for "How to troubleshoot LACP port channel groups".

### Syslog Message

ifgrp %s, port %s has transitioned to an inactive state. The interface group is in a degraded state.

### Parameters

**aggregateName** (STRING): Name of the interface group.

**portName** (STRING): Name of the member port.

## net.ipsconfig events

### net.ipsconfig.failed

#### Severity

ALERT

#### Description

This message occurs when creation of an IPspace fails. This failure leaves the system with an unusable IPspace.

#### Corrective Action

Delete the IPspace reporting the failure by using the "ipspace delete" command. Use the (privilege: diag) "debug smdb table ipspaces\_byname show" command to find the mapping from the ID reported in the event to the IPspace name required to delete this IPspace.

#### Syslog Message

IPspace (%d) needs to be deleted because configuration (%s) failed with error (%s).

#### Parameters

**ipsid** (INT): IPspace ID.

**configuration** (STRING): Configuration attempted.

**error** (STRING): Error returned.

## net.static events

### net.static.hostroute.replaced

#### Severity

NOTICE

#### Description

This event is generated when user configured static host route entry being replaced with a dynamic entry by route redirect message.

#### Corrective Action

(None).

## Syslog Message

rtredirect\_fib: User static route to destination(%s) through gateway(%s) being redirected to new gateway(%s)

## Parameters

**Dest** (STRING): IP address of the destination.

**OldGW** (STRING): IP address of the old gateway.

**NewGW** (STRING): IP address of the new gateway.

# net.vserverupdatefailure events

## net.vServerUpdateFailure

### Severity

NOTICE

### Description

This message occurs when a node cannot be updated with Vserver configuration. This can result in that node not being able to host LIFs.

### Corrective Action

(None).

### Syslog Message

Unable to update node %s with current Vservers.

### Parameters

**remote\_node** (STRING): Node being updated.

# net.vserverupdatesuccess events

## net.vServerUpdateSuccess

### Severity

NOTICE

### Description

This message occurs when the Vserver configuration on a node is successfully updated after a previous failure.

### Corrective Action

(None).

### Syslog Message

Vserver configuration is successfully updated on node %s.

### Parameters

**remote\_node** (STRING): Node that was updated.

# netgroup events

## netgroup.files events

### netgroup.files.missing

#### Severity

ERROR

#### Description

This message occurs when a netgroup lookup request finds that files is specified as a ns-switch source, but a netgroup file cannot be found.

#### Corrective Action

Check that the local netgroup file is present and load it if necessary. The commands to perform this are "vserver services netgroup load" and "vserver services netgroup status".

#### Syslog Message

Netgroup file missing for Vserver "%s" during lookup for netgroup "%s".

#### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

**netgroup** (STRING): Name of the netgroup for the lookup request.

## netgroup.ldap events

### netgroup.ldap.byhost.missing

#### Severity

INFORMATIONAL

#### Description

This message occurs when netgroup.byhost is disabled in the Lightweight Directory Access Protocol (LDAP) client configuration on the storage system, and LDAP is configured as an ns-switch source for the Vserver. Enabling netgroup.byhost enables mount operations to succeed faster when the netgroup size is large.

#### Corrective Action

Consider configuring the netgroup.byhost database on the LDAP server and enabling netgroup.byhost on the LDAP client configuration on the storage system by using the "ldap client modify" command.

#### Syslog Message

LDAP client config has netgroup.byhost disabled for Vserver "%s" during lookup for netgroup "%s".

#### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

**netgroup** (STRING): Name of the netgroup for the lookup request.

## netgroup.ldap.config

### Severity

ERROR

### Description

This message occurs when a netgroup lookup request finds that Lightweight Directory Access Protocol (LDAP) is specified as a ns-switch source, but LDAP is not configured for the Vserver. Netgroup lookups using LDAP will not function.

### Corrective Action

Check the ns-switch sources configured for the netgroup database using "vserver services name-service ns-switch show" and the LDAP configuration for the Vserver using "ldap show" and "ldap client show". Either remove LDAP as a ns-switch source, or configure LDAP.

### Syslog Message

LDAP configuration issues found for Vserver "%s" during lookup for netgroup "%s".

### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.  
**netgroup** (STRING): Name of the netgroup for the lookup request.

## netgroup.nis events

### netgroup.nis.byhost.decode

#### Severity

ERROR

#### Description

This message occurs when the netgroup.byhost remote procedure call to the Network Information Service (NIS) server returns a response that cannot be decoded.

#### Corrective Action

Check the netgroup byhost database on the NIS server for lines that have more than 1024 characters.

#### Syslog Message

NIS server returned an undecodable response for Vserver "%s" during lookup for netgroup "%s".

#### Parameters

**vserver** (STRING): Name of the Vserver on which the problem is occurring.  
**netgroup** (STRING): Name of the netgroup for the lookup request.

### netgroup.nis.byhost.missing

#### Severity

INFORMATIONAL

#### Description

This message occurs when the netgroup.byhost map is not configured on the Network Information Service (NIS) server and NIS is configured as a ns-switch source for the Vserver. Enabling netgroup.byhost enables

mount operations to succeed faster when the netgroup size is large.

### **Corrective Action**

Consider configuring the netgroup.byhost map on the NIS server to gain better performance with netgroups.

### **Syslog Message**

NIS server missing netgroup.byhost map for Vserver "%s" during lookup for netgroup "%s".

### **Parameters**

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

**netgroup** (STRING): Name of the netgroup for the lookup request.

## **netgroup.nis.config**

### **Severity**

ERROR

### **Description**

This message occurs when a netgroup lookup request finds that Network Information Service (NIS) is specified as a ns-switch source, but NIS is not configured for the Vserver. Netgroup lookups using NIS will not function.

### **Corrective Action**

Check the ns-switch sources configured for the netgroup database using "vserver services name-service ns-switch show" and the NIS configuration for the Vserver using "nis-domain show". Either remove NIS as a ns-switch source, or configure NIS.

### **Syslog Message**

NIS configuration issues found for Vserver "%s" during lookup for netgroup "%s".

### **Parameters**

**vserver** (STRING): Name of the Vserver on which the problem is occurring.

**netgroup** (STRING): Name of the netgroup for the lookup request.

# netif events

## netif.autonegfailed events

### netif.autoNegFailed

#### Severity

ERROR

#### Description

This message occurs when autonegotiation fails and the interface is operating at half-duplex, which might not be the desired mode.

#### Corrective Action

Verify that autonegotiation is enabled on the switch or device to which the interface is connected. Or, if the desired speed is 10 or 100 Mbps, disable autonegotiation on the interface by setting the mediatype.

#### Syslog Message

Autonegotiation failed, interface %s operating at %d Mbps half-duplex.

#### Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

**speed** (INT): Speed used by the interface.

## netif.badeeprom events

### netif.badEeprom

#### Severity

EMERGENCY

#### Description

This message occurs when the network interface has an electronically erasable programmable read-only memory (EEPROM) with a bad cyclic redundancy check (CRC) or media access control (MAC) address.

#### Corrective Action

Replace the component associated with the network interface (PCIe card or controller motherboard).

#### Syslog Message

Network interface %s has an EEPROM with a bad CRC or MAC address.

#### Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

## netif.badoversizepacket events

## netif.badOversizePacket

### Severity

ERROR

### Description

This message occurs when the network interface card (NIC) receives a bad oversized packet.

### Corrective Action

Try a different switch port.

### Syslog Message

Bad oversized packet received on network interface %s. Check for NIC receive errors (CRC, runts, fragments). Might indicate a bad switch port.

### Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

## netif.cpreventcompleted events

### netif.cprEventCompleted

### Severity

INFORMATIONAL

### Description

This message occurs when a CPR Event (core PCI resiliency) has completed.

### Corrective Action

PCI reset has occurred automatically to correct.

### Syslog Message

CPR Event (core PCI resiliency) completed on network interface %s and the count since boot is %d.

### Parameters

**ifname** (STRING): Name of network interface.

**count** (INT): Number of times a CPR Event (core PCI resiliency) has been completed since boot.

## netif.cpreventedected events

### netif.cprEventDetected

### Severity

INFORMATIONAL

### Description

This message occurs when a CPR Event (core PCI resiliency) is detected.



### Corrective Action

PCI reset will occur automatically to correct.

### Syslog Message

CPR Event (core PCI resiliency) detected on network interface %s and the count since boot is %d.

### Parameters

**ifname** (STRING): Name of network interface.

**count** (INT): Number of times CPR Event (core PCI resiliency) has occurred since boot.

## netif.cprmaxeventsreached events

### netif.cprMaxEventsReached

#### Severity

ERROR

#### Description

This message occurs when the maximum number of CPR Events (core PCI resiliency) has occurred since boot.

#### Corrective Action

Replace the failing hardware.

#### Syslog Message

Replace hardware — maximum number of CPR Events (core PCI resiliency) since boot has occurred on network interface %s.

#### Parameters

**ifname** (STRING): Name of network interface.

## netif.eepromreadfailed events

### netif.eepromReadFailed

#### Severity

ERROR

#### Description

This message occurs when there is a failure in reading the electronically erasable programmable read-only memory (EEPROM) for a network interface. The interface is not initialized.

#### Corrective Action

Rebooting the system might clear the problem. The problem might be caused by a software bug and should be reported to NetApp technical support.

#### Syslog Message

Failure in reading EEPROM for network interface %s.

## Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

# netif.excesslinkinterrupts events

## netif.excessLinkInterrupts

### Severity

ERROR

### Description

This message occurs when excessive link status change interrupts occur on the network interface. One or two events might indicate a momentary glitch on the link, but many events might indicate a bad switch port, cable, or network interface card (NIC).

### Corrective Action

Try using a different switch port and cable. If the problem persists, replace the NIC.

### Syslog Message

Excessive link status change interrupts occurred on network interface %s.

## Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

# netif.fatal events

## netif.fatal.err

### Severity

ALERT

### Description

This message occurs when the network device's driver encounters an error from which it cannot recover.

### Corrective Action

Replace the network interface card.

### Syslog Message

The network device in slot %d encountered fatal error %s.

## Parameters

**slot** (INT): Slot number.

**info** (STRING): Device specific information

# netif.fwupdatefailed events

## netif.fwUpdateFailed

**Severity**

ERROR

**Description**

This message occurs when there is an error during a network interface firmware update.

**Corrective Action**

In very rare cases, a firmware update failure might leave the interface inoperable. If the driver initialization routine subsequently fails and logs a message, the motherboard or card might need to be replaced. For more information or assistance, contact NetApp technical support.

**Syslog Message**

Firmware update of network interface %s to version %s failed :%s.

**Parameters**

**ifName** (STRING): Name of the network interface.

**version** (STRING): Version of the firmware update.

**reason** (STRING): Reason for the firmware update failure.

## netif.fwupdatestatus events

### netif.fwUpdateStatus

**Severity**

NOTICE

**Description**

This message occurs to indicate the status of a network interface firmware update.

**Corrective Action**

(None).

**Syslog Message**

Firmware update of network interface %s to version %s %s.

**Parameters**

**ifName** (STRING): Name of the network interface.

**version** (STRING): Version of firmware update.

**status** (STRING): Status of firmware update process (started or completed)

## netif.hangdetected events

### netif.hangDetected

**Severity**

ERROR

**Description**

This message occurs when a network interface driver detects that an interface is hung.

## Corrective Action

If hardware is reset and there are no further hangs, no corrective action is required. Otherwise, if the system reports frequent hangs, a core dump is forced, or no action is taken by the driver, contact NetApp technical support.

## Syslog Message

Network interface %s hung (%s). %s. Driver: %s.

## Parameters

**ifName** (STRING): Name of the network interface: for example, "e4a".

**hangType** (STRING): Type of hang. For example, "transmit", "PCIe RcvMstAbt".

**recoveryAction** (STRING): Action, if any, that will be taken by the driver to recover. For example, "Resetting to recover" or "Forcing panic for core dump" (for root cause analysis).

**driver** (STRING): Driver name.

# netif.hn events

## netif.hn.setDataPathFailure

### Severity

ERROR

### Description

This message occurs when Hyper-V virtual NIC fails to switch the data path from a Physical Function (PF) vPort to a Virtual Function (VF) vPort or vice-versa when accelerated networking is enabled on the virtual NIC. When the data path switch fails, the traffic might or might not continue to flow through existing path. If no traffic flows, then it can result in a complete outage on the LIFs hosted by the virtual NIC.

### Corrective Action

Redeploy the CVO instance on another host machine. If the problem persists, contact NetApp technical support.

### Syslog Message

%s: Failed to set data path to %s. Error: %d.

### Parameters

**ifName** (STRING): Name of the network interface.

**path** (STRING): Data path: synthetic (via PF) or virtual (via VF).

**error** (INT): Failure to set data path.

# netif.init events

## netif.init.failed

### Severity

ALERT

### Description

This message occurs when initialization of a network interface failed due to an unexpected software error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Initialization of network interface %s failed due to unexpected software error %s:%d.

### Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

**module** (STRING): Name of the software module or driver.

**error** (INT): Module-specific error code.

## netif.linkdown events

### netif.linkDown

#### Severity

INFORMATIONAL

#### Description

This message occurs when a network link goes down or the interface is not operational when first configured up.

#### Corrective Action

If the message describes the link as "down," check the cabling and switch port configuration. If the message states "controller initialization failure", report this to NetApp technical support.

#### Syslog Message

%s %s: Link %s.

#### Parameters

**lanType** (STRING): Local area network (LAN) type.

**ifName** (STRING): Name of the network interface.

**status** (STRING): Brief information about the problem.

## netif.linkerrors events

### netif.linkErrors

#### Severity

ERROR

#### Description

This message occurs when the driver detects an excessive link error rate. Link errors are cyclic redundancy checks (CRC), runt frames, fragment, jabber, and alignment errors. Use the 'ifstat' command to display counts.

#### Corrective Action

Verify that the cable is securely connected at both ends. Verify that both devices are configured the same with respect to autonegotiation, speed, and duplex. If the problem persists, there might be bad hardware at one end, or a bad cable or connector.

## Syslog Message

Excessive link errors on network interface %s. Might indicate a bad cable, switch port, or NIC, or that a cable connector is not fully inserted in a socket. On a 10/100 port, might indicate a duplex mismatch.

## Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

# netif.linkinfo events

## netif.linkInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a network interface is reset by the software or is configured down by the administrator.

### Corrective Action

(None).

## Syslog Message

%s %s: Link %s.

## Parameters

**lanType** (STRING): Local area network (LAN) type.

**ifName** (STRING): Name of the network interface.

**info** (STRING): Brief description of the event.

# netif.linkup events

## netif.linkUp

### Severity

INFORMATIONAL

### Description

This message occurs when the state of a network link changes to "up".

### Corrective Action

(None).

## Syslog Message

%s %s: Link %s.

## Parameters

**lanType** (STRING): Local area network (LAN) type.

**ifName** (STRING): Name of the network interface.

**status** (STRING): Status of the link, "up".

## netif.nicnotsupported events

### netif.nicNotSupported

#### Severity

ERROR

#### Description

This message occurs when the installed network interface card (NIC) version is not supported.

#### Corrective Action

Replace the NIC with one purchased from NetApp.

#### Syslog Message

The NIC in slot %d is not supported. The %s is not supported.

#### Parameters

**slot** (INT): Slot number of the NIC.

**reason** (STRING): Specific reason the NIC is not supported.

## netif.overtemperror events

### netif.overTempError

#### Severity

EMERGENCY

#### Description

This message occurs when a network port is too hot or a sensor has malfunctioned. Depending on how the port is deployed, the failure may disrupt user data access, node access or cluster communication.

#### Corrective Action

Ensure adequate cooling and ventilation is available. If that does not resolve the issue, replace the failed card.

#### Syslog Message

Network port %s is over temperature.

#### Parameters

**ifName** (STRING): Name of the network port: for example, "e4a".

## netif.phynotread events

### netif.phyNotRead

#### Severity

ERROR

## Description

This message occurs when the system cannot read identification information from the physical layer (PHY) of a network interface. This might be caused by a software bug or hung hardware, and the link might not come up.

## Corrective Action

If the link does not come up, power-cycle (not just reboot) the system. If the link still does not come up, replace the motherboard or network interface card, depending on the location of the interface.

## Syslog Message

Cannot read PHY for network interface %s, status mask=0x%llx.

## Parameters

**ifName** (STRING): Name of the network interface, for example, "e8a".

**statusMask** (LONGINTEX): Technical information about the event for Data ONTAP® developers.

# netif.phyoperationfailed events

## netif.phyOperationFailed

### Severity

ERROR

### Description

This message occurs when there is a failure in reading or writing the physical layer (PHY) of a network interface.

### Corrective Action

This might be caused by a software bug. If there are problems with the interface after seeing this message, replace the component associated with the network interface (PCIe card or controller motherboard). If the problem persists, contact NetApp technical support.

### Syslog Message

Failure in reading or writing PHY for network interface %s. Driver=%s, error=%d, regaddr=0x%x, command=0x%x, status=0x%x.

### Parameters

**ifName** (STRING): Name of the network interface; for example, "e2a".

**driver** (STRING): Driver name.

**error** (INT): Error code.

**regaddr** (INTEX): PHY register address.

**command** (INTEX): PHY command.

**status** (INTEX): PHY status.

# netif.phyread events

## netif.phyRead

### Severity

ERROR



## Description

This message occurs when there is a failure in reading the physical layer (PHY) of a network interface.

## Corrective Action

This might be caused by a software bug. If there are problems with the interface after seeing this message, contact NetApp technical support.

## Syslog Message

Failure in reading PHY for network interface %s.

## Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

# netif.ratelimitthreshold events

## netif.rateLimitThreshold

### Severity

ERROR

### Description

This message occurs when the protocol rate threshold is reached on a network interface.

### Corrective Action

Fix the faulty network configuration or incorrect setup that enables a sudden spike in broadcast packets to bring down the node. If you still want a high ARP rate threshold, use the "bootarg.arp.ratelimit.threshold" boot argument to set that threshold.

### Syslog Message

High rate limit on network interface %s for broadcast protocol %s being detected: %d pkts/sec.

### Parameters

**ifName** (STRING): Name of the network interface that reached the protocol rate threshold.

**protocol** (STRING): Name of the broadcast protocol that overloaded the network.

**packetRate** (INT): Observed incoming packet rate per second on the interface.

# netif.sfpeventerrorcode events

## netif.sfpEventErrorCode

### Severity

ERROR

### Description

This message occurs when the installed Small Form-factor Pluggable (SFP+ or QSFP+) module is either not supported or faulty.

### Corrective Action

The cable could be unsupported or faulty. Replace the cable and check the switch port configuration for errors.

## Syslog Message

Unsupported or faulty transceiver or cable in port %s. Error :%s.

## Parameters

**ifName** (STRING): Name of the network interface.

**error\_string** (STRING): module event error string.

# netif.sfpnotsupported events

## netif.sfpNotSupported

### Severity

ERROR

### Description

This message occurs when the installed Small Form-factor Pluggable (SFP+ or QSFP+) module is not supported.

### Corrective Action

Replace the SFP+ or QSFP+ module with one purchased from NetApp for this interface.

### Syslog Message

The SFP+ or QSFP+ module (%s %s) installed in %s is not supported with this network interface.

### Parameters

**sfpVendor** (STRING): SFP+ or QSFP+ vendor name.

**sfpPn** (STRING): SFP+ or QSFP+ part number.

**ifName** (STRING): Name of the network interface; for example, "e2a".

# netif.speeddowngraded events

## netif.speedDowngraded

### Severity

ERROR

### Description

This message occurs when the interface is operating at a lower speed than expected.

### Corrective Action

The physical connection may be faulty. Replace the cable and check the switch port configuration for errors.

### Syslog Message

Interface %s operating at %d Mbps. This might occur due to faulty cable, faulty switch hardware, or faulty switch software.

### Parameters

**ifName** (STRING): Name of the network interface.

**speed** (INT): Speed used by the interface.

# netif.speednotsupported events

## netif.speedNotSupported

### Severity

ERROR

### Description

This message occurs when a network interface autonegotiates to an unsupported speed.

### Corrective Action

Connect to 1 or 10 gigabits per second (Gbps) switch port.

### Syslog Message

Autonegotiation on network interface %s resolved to speed %d. This speed is not supported.

### Parameters

**ifname** (STRING): Name of network interface.

**speed** (INT): Speed in megabits per second (Mbps).

# netif.tcp events

## netif.tcp.conn.bad.checksum

### Severity

ERROR

### Description

This message occurs when a bad checksum is detected in the incoming TCP packets.

### Corrective Action

If experiencing network issues, corrupted TCP packets might be an issue. Check the network components along the path of TCP connection.

### Syslog Message

TCP packet with bad checksum detected on port %s. The packet arrived on connection with source address %s and destination\_address %s.

### Parameters

**local\_port** (STRING): Ethernet port where the TCP bad checksum packets were detected.

**source\_address** (STRING): Source IP address of the connection.

**destination\_address** (STRING): Destination IP address of the connection.

# netif.toomanynics events

## netif.tooManyNics

### Severity

ERROR

## Description

This message occurs when a network interface was not initialized because the number of interfaces of this type exceeds the maximum supported.

## Corrective Action

Remove a network interface card (NIC).

## Syslog Message

Network interface %s was not initialized because the number of interfaces of this type exceeds the maximum (%d) supported.

## Parameters

**ifName** (STRING): Name of the network interface; for example, "e8a".

**maxIf** (INT): Maximum number of interfaces of this type that is supported.

# netif.uncoreccerror events

## netif.uncorEccError

### Severity

EMERGENCY

### Description

This message occurs when the network interface controller reports an unrecoverable ECC error.

### Corrective Action

Reboot the controller. If the problem persists, replace the component associated with the network interface (PCIe card or controller motherboard). If the problem persists, contact NetApp technical support.

### Syslog Message

Unrecoverable ECC error on network interface %s.

### Parameters

**ifName** (STRING): Name of the network interface; for example, "e4a".

# netif.unknownswrequest events

## netif.unknownSwRequest

### Severity

ERROR

### Description

This message occurs when a network driver does not recognize a request from another software module, possibly because the request is invalid.

### Corrective Action

Power-cycle the controller; a reboot is not sufficient. If the problem persists, contact NetApp technical support.

## Syslog Message

Software module %s was passed unknown %s %d for %s.

## Parameters

**driver** (STRING): Driver module reporting the error.

**object** (STRING): Name of the requested object.

**value** (INT): Value that is not recognized.

**ifName** (STRING): Name of the network interface.

# netinet events

## netinet.ethr events

### netinet.ethr.arpstorm

#### Severity

ERROR

#### Description

This message occurs when a high arrival rate of address resolution protocol (ARP) frames has been detected (possible ARP storm). This condition can indicate a possible network configuration issue.

#### Corrective Action

If you are experiencing network issues, an ARP storm might be a concern. Confirm the ARP storm event is valid and check the network configuration.

#### Syslog Message

The node detected a possible ARP storm on port %s. Detected rate: %d packets per second.

#### Parameters

**port\_name** (STRING): Port where the possible ARP storm was detected.

**rate\_count** (LONGINT): ARP arrival rate seen at the port where the possible ARP storm was detected.

### netinet.ethr.broadcastAddr

#### Severity

ERROR

#### Description

This message occurs when the system finds that the Ethernet address for the IP address is found to be a broadcast address.

#### Corrective Action

(None).

#### Syslog Message

ARP: Ethernet address is a broadcast address for IP address %s!

#### Parameters

**ipAddress** (STRING): IP address in Standard Internet dotted notation.

### netinet.ethr.dup.clustIP

#### Severity

EMERGENCY

#### Description

This message occurs when a new node that is trying to join the cluster is using the same IP address that is configured on this node's cluster interface. This can prevent data from being served to clients and

potentially bring down the entire cluster.

### Corrective Action

Use the Ethernet (MAC) address with the "network port show -mac <mac>" command to find the network port. Use the "network interface show -curr-port <port>" command to identify the cluster interface that is configured with the duplicate IP address. Remove the LIF on the cluster interface of the new node in the cluster network by using the "net int delete -vserver <vserver\_name> -lif <lif\_name>" command. After removing the LIF, assign a new, unique IP address to the cluster interface.

### Syslog Message

Duplicate IP address %s (Ethernet MAC: %s) configured on the cluster network.

### Parameters

**ipaddr** (STRING): Duplicate IP address that was detected on the cluster network.

**ethaddr** (STRING): Ethernet (MAC) address that uniquely identifies the node configured with the duplicate IP address.

## netinet.ethr.duplct.ipAdrr

### Severity

ERROR

### Description

This message occurs when another system is detected using an IP address configured on this system.

### Corrective Action

The Ethernet address uniquely identifies the other system that is using the IP address. If you are using static IP addresses, change the IP address configured on one of the systems. If DHCP is in use on the network, ensure that the static IP address is not in the range being used by the DHCP server.

### Syslog Message

Duplicate IP address %s!! sent from Ethernet address: %s.

### Parameters

**ipaddr** (STRING): IP address.

**ethaddr** (STRING): Ethernet address.

## netinet.icmp events

### netinet.icmp.redirect

### Severity

NOTICE

### Description

This message occurs when an ICMPv4 redirect message is received. The redirect might be received due to either a misconfigured route or a man-in-the-middle ICMP redirect attack.

### Corrective Action

When enabled, ICMP redirects can be disabled using the "network tuning icmp modify -is-drop-redirect -enabled true" command. However, disabling ICMP redirects might lead to connectivity problems or longer

delays of data transfers.

### **Syslog Message**

Received %lu ICMPv4 redirect messages for IPspace "%s".

### **Parameters**

**no\_of\_redirects** (LONGINT): Total number of ICMPv4 redirects received for the specific IPspace.  
**ipspace** (STRING): The IPspace for which the redirect packets are received.



# netinet6 events

## netinet6.frag events

### netinet6.frag.bytes.overlap

**Severity**

ERROR

**Description**

This message occurs when the incoming fragment overlaps some of the existing previous fragment in the reassembly queue.

**Corrective Action**

(None).

**Syslog Message**

%d bytes of a fragment from %s overlapped the %s fragment

**Parameters**

**bytes** (INT): Number of bytes.

**ipAddress** (STRING): Printable that is loggable representation of IP address.

**message** (STRING): String passed as "previous" or "succeeding".

## netinet6.icmp events

### netinet6.icmp.chksum

**Severity**

ERROR

**Description**

This message occurs when the system receives an Internet Control Message Protocol(ICMP) message, and while processing the ICMP message it finds a checksum error.

**Corrective Action**

(None).

**Syslog Message**

ICMP6 checksum error (%d|%x) %s.

**Parameters**

**icmpType** (INT): ICMP header type field.

**checksum** (INTHEX): Checksum value.

**ipAddress** (STRING): Printable that is a loggable representation of IP address.

### netinet6.icmp.redirect

**Severity**

NOTICE

**Description**

This message occurs when an ICMPv6 redirect message is received. The redirect might be received due to either a misconfigured route or a man-in-the-middle ICMP redirect attack.

**Corrective Action**

When enabled, ICMP redirects can be disabled using the "network tuning icmp6 modify -is-v6-redirect -accepted false" command. However, disabling ICMP redirects might lead to connectivity problems or longer delays of data transfers.

**Syslog Message**

Received %lu ICMPv6 redirect messages for IPspace "%s".

**Parameters**

**no\_of\_redirects** (LONGINT): Total number of ICMPv6 redirects received for the specific IPspace.  
**ipspace** (STRING): The IPspace for which the redirect packets are received.

**netinet6.icmp.reject.equal****Severity**

ERROR

**Description**

This message occurs when the system rejects a redirect because both router case and on-link destination case are equal.

**Corrective Action**

(None).

**Syslog Message**

ICMP6 redirect rejected; neither router case nor onlink case: %s.

**Parameters**

**addressBuff** (STRING): Contains IP6 addresses.

**netinet6.icmp.reject.length****Severity**

INFORMATIONAL

**Description**

This message occurs when there is a lladdrlen mismatch for the specified IP6 address.

**Corrective Action**

(None).

**Syslog Message**

icmp6\_redirect\_input: lladdrlen mismatch for %s (if %d, icmp6 packet %d): %s.

## Parameters

**ip6address** (STRING): Printable that is a loggable representation of the IP address.  
**mediaAddrLen** (INT): Media address length.  
**lladdrLen** (INT): Address length.  
**addressBuff** (STRING): Contains IP6 addresses.

## netinet6.icmp.reject.link

### Severity

ERROR

### Description

This message occurs when the system rejects an Internet Control Message Protocol6 (ICMP6) redirect.

### Corrective Action

(None).

### Syslog Message

ICMP6 redirect sent from %s rejected;It must be from the linklocal address.

## Parameters

**ipAddress** (STRING): Printable that is a loggable representation of the IP address.

## netinet6.icmp.reject.ndopt

### Severity

INFORMATIONAL

### Description

This message occurs because of invalid neighbor discovery options.

### Corrective Action

(None).

### Syslog Message

icmp6\_redirect\_input: Invalid neighbor discovery option, rejected: %s.

## Parameters

**addressBuff** (STRING): Contains IP6 addresses.

## netinet6.in6 events

### netinet6.in6.invlid.prfxLen1

### Severity

ERROR

### Description

This message occurs because of equal prefix length.

**Corrective Action**

(None).

**Syslog Message**

in6\_are\_prefix\_equal: Invalid prefix length (%d).

**Parameters**

**prfxLen** (INT): Prefix length.

**netinet6.in6.invlid.prfxLen2****Severity**

ERROR

**Description**

This message occurs because of invalid prefix length.

**Corrective Action**

(None).

**Syslog Message**

in6\_prefixlen2mask: Invalid prefix length (%d).

**Parameters**

**prfxLen** (INT): Prefix length.

**netinet6.input events****netinet6.input.addr.notReady****Severity**

INFORMATIONAL

**Description**

This message occurs when the destination address is not ready, so the packet is discarded.

**Corrective Action**

(None).

**Syslog Message**

ip6\_input: packet was sent to an unready address %s→%s.

**Parameters**

**srcAddr** (STRING): Source IP6 address.

**dstAddr** (STRING): Destination IP6 address.

**netinet6.input.opt.payload****Severity**

ERROR

**Description**

This message occurs when the length of the payload option in the header is inconsistent.

**Corrective Action**

(None).

**Syslog Message**

Length of jumbopayload option is inconsistent(%d).

**Parameters**

**option** (INT): Length of option.

**netinet6.input.opt.rtrAlert****Severity**

ERROR

**Description**

This message occurs when the length of the router alert option in the header is inconsistent.

**Corrective Action**

(None).

**Syslog Message**

Length of router alert option is inconsistent(%d).

**Parameters**

**option** (INT): Length of option.

**netinet6.mld6 events****netinet6.mld6.illegal.type****Severity**

ERROR

**Description**

This message occurs when the system discovers an illegal Internet Control Message Protocol (ICMP) type field in a multicast listener discovery structure.

**Corrective Action**

(None).

**Syslog Message**

mld6\_input: Illegal ICMP type field (%d).

**Parameters**

**type** (INT): ICMP type field.

## netinet6.mld6.srcAddr.invl

### Severity

ERROR

### Description

This message occurs during validation of the source address, when the system finds that the source address is a non-linked local address.

### Corrective Action

(None).

### Syslog Message

mld6\_input: Source address %s is not link-local.

### Parameters

**srcAddr** (STRING): Source IP6 address.

## netinet6.nbr events

### netinet6.nbr.bad.addr

### Severity

ERROR

### Description

This message occurs when ONTAP receives unexpected Neighbor Discovery (ND) packets. The packets could come from various sources, but usually they come from a source IP address. The EMS message shows the information from the packet. If something specific is wrong, another EMS message with more information will display with this EMS message.

### Corrective Action

ONTAP received an invalid Neighbor Discovery (ND) packet. If the condition persists, check your network infrastructure, such as switches, routers, or other connected equipment.

### Syslog Message

nd6\_ns\_input: src=%s, dst=%s, tgt=%s.

### Parameters

**srcIpAddress** (STRING): Source IPv6 address.

**dstIpAddress** (STRING): Destination IPv6 address.

**tgtIpAddress** (STRING): Target IPv6 address.

### netinet6.nbr.bad.tgtAddr

### Severity

INFORMATIONAL

### Description

This message occurs when the target address packet is found to be multicast.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: Bad NS target (multicast).

**Parameters**

(None).

**netinet6.nbr.badDstAddr.pkt1****Severity**

INFORMATIONAL

**Description**

This message occurs when the destination address packet is bad due to a wrong destination IP6 address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: Bad DAD packet (wrong ip6 dst).

**Parameters**

(None).

**netinet6.nbr.badDstAddr.pkt2****Severity**

INFORMATIONAL

**Description**

This message occurs when the destination address packet is bad due to link-layer address option.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: Bad DAD packet (link-layer address option).

**Parameters**

(None).

**netinet6.nbr.dad.complete****Severity**

ALERT

**Description**

This message occurs when Duplicate address Detection (DAD) for the specified interface address is completed, and duplicate address is found.

## Corrective Action

(None).

## Syslog Message

%s: DAD complete for %s - duplicate found.

## Parameters

**ifAddress** (STRING): Interface address.

**ip6address** (STRING): Printable that is a loggable representation of the IP address.

## netinet6.nbr.dad.ignrNS

### Severity

INFORMATIONAL

### Description

This message occurs when Duplicate address Detection (DAD) ignores Neighbor Solicitation (NS) packet for the specified interface address.

### Corrective Action

(None).

### Syslog Message

nd6\_dad\_ns\_input: Ignoring DAD NS packet for address %s(%s).

### Parameters

**targetAddress** (STRING): Printable that is loggable representation of target IPv6 address.

**ifAddress** (STRING): Interface address.

## netinet6.nbr.dad.memAlcFail

### Severity

ALERT

### Description

This message occurs when Duplicated Address Detection (DAD) is started and memory allocated fails for the specified interface address.

### Corrective Action

(None).

### Syslog Message

nd6\_dad\_start: Memory allocation failed for %s(%s).

### Parameters

**ip6address** (STRING): Printable that is a loggable representation of an IP6 address.

**ifAddress** (STRING): Interface address.

## netinet6.nbr.dad.timeout



**Severity**

NOTICE

**Description**

This message occurs during Duplicated Address Detection (DAD), when number of tries or count to transmit DAD packet exceeds the predefined maximum number of tries.

**Corrective Action**

(None).

**Syslog Message**

%s: could not run DAD, driver or link problem?

**Parameters**

**ifAddress** (STRING): Interface address.

**netinet6.nbr.dadtmr.duplcAdr****Severity**

NOTICE

**Description**

This message occurs when Duplicated Address Detection (DAD) timer is called, and the parameter passed to it, is a duplicated interface address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_dad\_timer: Called with duplicated address %s(%s).

**Parameters**

**ip6address** (STRING): Printable that is a loggable representation of an IP6 address.

**ifAddress** (STRING): Interface address.

**netinet6.nbr.dadTmr.nTntvAdr****Severity**

NOTICE

**Description**

This message occurs when Duplicated Address Detection (DAD) Timer is called with non tentative interface address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_dad\_timer: Called with non-tentative address %s(%s).

## Parameters

**ip6address** (STRING): Printable that is a loggable representation of an IP6 address.

**ifAddress** (STRING): Interface address.

## netinet6.nbr.dadtmr.nullPrm

### Deprecated

Deprecated as of Data ONTAP 9.8 since this EMS is no longer needed.

### Severity

NOTICE

### Description

This message occurs when Duplicated Address Detection (DAD) timer is called and the interface address passed to it is NULL.

### Corrective Action

(None).

### Syslog Message

nd6\_dad\_timer: Called with null parameter.

### Parameters

(None).

## netinet6.nbr.duplcte.taddr

### Severity

ALERT

### Description

This message occurs when there is a duplicate target IP6 address.

### Corrective Action

(None).

### Syslog Message

nd6\_na\_input: Duplicate IP6 address %s.

### Parameters

**targetAddress** (STRING): Printable that is a loggable representation of target IP6 address.

## netinet6.nbr.duplcteSrc.addr

### Severity

INFORMATIONAL

### Description

This message occurs when there is duplicate source IP6 address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: Duplicate IP6 address %s.

**Parameters**

**srcAddress** (STRING): Printable that is a loggable representation of source IP6 address.

**netinet6.nbr.invlD.ndOpt1****Severity**

INFORMATIONAL

**Description**

This message occurs because of invalid Neighbor Discovery (ND)option.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: Invalid ND option; ignored.

**Parameters**

(None).

**netinet6.nbr.invlD.ndOpt2****Severity**

INFORMATIONAL

**Description**

This message occurs because of an invalid Neighbour Discovery (ND) option during neighbor advertising input handling.

**Corrective Action**

(None).

**Syslog Message**

nd6\_na\_input: Invalid ND option; ignored.

**Parameters**

(None).

**netinet6.nbr.invlD.tgtAddr****Severity**

NOTICE

**Description**

This message occurs during Neighbor Advertising (NA)input handling, because of invalid target address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_na\_input: Invalid target address %s.

**Parameters**

**targetAddress** (STRING): Target IP6 address.

**netinet6.nbr.manl.intvtnReq****Severity**

ERROR

**Description**

This message occurs when 'dad\_duplicated' is called and manual intervention is required.

**Corrective Action**

(None).

**Syslog Message**

%s: Manual intervention required.

**Parameters**

**ifAddress** (STRING): Interface address.

**netinet6.nbr.misMth.lladrIn1****Severity**

INFORMATIONAL

**Description**

This message occurs when there is a lladrIn mismatch for the specified IP6 address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ns\_input: lladrIn mismatch for %s (if %d, NS packet %d).

**Parameters**

**ip6address** (STRING): Printable that is a loggable representation of an IP address.

**ifAddrIn** (INT): Interface address length.

**lladrIn** (INT): Address length.

**netinet6.nbr.misMth.lladrIn2**

**Severity**

INFORMATIONAL

**Description**

This message occurs during Neighbor Advertising (NA) input handling, when there is a lladdrLen mismatch for the specified IP6 address.

**Corrective Action**

(None).

**Syslog Message**

nd6\_na\_input: lladdrLen mismatch for %s (if %d, NA packet %d).

**Parameters**

**ip6address** (STRING): Printable that is a loggable representation of target IP address.

**ifAddrLen** (INT): Interface address length.

**lladdrLen** (INT): Address length.

## netinet6.nbr.solAdv.mlticstd

**Severity**

NOTICE

**Description**

This message occurs during Neighbor Advertising(NA) input handling, as solicited advertising is multicast.

**Corrective Action**

(None).

**Syslog Message**

nd6\_na\_input: A solicited advertising is multicast.

**Parameters**

(None).

## netinet6.nd6 events

### netinet6.nd6.setmtu.small

**Severity**

NOTICE

**Description**

This message occurs when the port MTU (Maximum Transmission Unit) is set to a smaller value than the IPv6 minimum MTU of 1280. If you set the port MTU value to less than 1280, which is not recommended, IPv6 packets larger than the set value cannot be sent from or received by the controller.

**Corrective Action**

Use the 'broadcast-domain modify' command with the '-mtu' option to increase MTU of all ports in the domain to at least the IPv6 minimum MTU of 1280.

## Syslog Message

nd6\_setmtu: new port MTU '%lu' on the port '%s' is too small for IPv6.

## Parameters

**linkmtu** (LONGINT): Configured MTU value for the port.

**ifName** (STRING): Name of the port on which the MTU is being set.

# netinet6.rtr events

## netinet6.rtr.high.mtu

### Severity

INFORMATIONAL

### Description

This message occurs when the Maximum Transmission Unit (MTU) size of the Router Advertisement (RA) message, sent from the specified source IPv6 address, is higher than the upper bound of the Link MTU.

### Corrective Action

(None).

## Syslog Message

nd6\_ra\_input: MTU option mtu=%d sent from %s; exceeds the maxmtu %d on '%s'; ignoring.

## Parameters

**mtu** (INT): MTU size.

**ip6Address** (STRING): Printable that is a loggable representation of the source IPv6 address.

**maxmtu** (INT): Upper bound of the Link MTU.

**ifName** (STRING): Interface name.

## netinet6.rtr.invlld.ndOpt1

### Severity

INFORMATIONAL

### Description

This message occurs because of an invalid Neighbor Discovery(ND) option.

### Corrective Action

(None).

## Syslog Message

nd6\_rs\_input: Invalid ND option; ignored.

## Parameters

(None).

## netinet6.rtr.invlld.ndOpt2

**Severity**

INFORMATIONAL

**Description**

This message occurs because of an invalid Neighbor Discovery(ND) option in the receiving Router Advertising (RA)message.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ra\_input: Invalid ND option; ignored.

**Parameters**

(None).

**netinet6.rtr.invlD.optLength****Severity**

INFORMATIONAL

**Description**

This message occurs because of an invalid option length for a prefix information option, when receiving a Router Advertising (RA)message. The option length must be equal to 4. This invalid option length is ignored.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ra\_input: Invalid option length %d for prefix information option; ignored.

**Parameters**

**optLength** (INT): Option length.

**netinet6.rtr.invlD.prefix****Severity**

INFORMATIONAL

**Description**

This message occurs because of an invalid prefix in the prefix information structure, when receiving Router Advertising(RA) message.

**Corrective Action**

(None).

**Syslog Message**

nd6\_ra\_input: Invalid prefix %s; ignored.

## Parameters

**prefix** (STRING): Prefix.

## netinet6.rtr.invlD.prfxLen1

### Severity

INFORMATIONAL

### Description

This message occurs because of an invalid prefix length for prefix information option, when receiving Router Advertising(RA) message. The prefix length must be less than 128.

### Corrective Action

(None).

### Syslog Message

nd6\_ra\_input: Invalid prefix length %d for prefix information option; ignored.

## Parameters

**prfxLength** (INT): Prefix length.

## netinet6.rtr.low.mtu

### Severity

INFORMATIONAL

### Description

This message occurs when the Maximum Transmission Unit (MTU) size is lower than 1,280, which is the minimal MTU and reassembly size.

### Corrective Action

(None).

### Syslog Message

nd6\_ra\_input: MTU option mtu=%d sent from %s is less than the minimum Link MTU on '%s'; ignoring.

## Parameters

**mtu** (INT): MTU size.

**ip6Address** (STRING): Printable that is a loggable representation of the source IPv6 address.

**ifName** (STRING): Interface name.

## netinet6.rtr.misMth.lldrn1

### Severity

INFORMATIONAL

### Description

This message occurs when there is a lladdrln mismatch for the specified IP6 address of the Router Solicitation (RS) packet.



## Corrective Action

(None).

## Syslog Message

nd6\_rs\_input: lladdrlen mismatch for %s (if %d, RS packet %d).

## Parameters

**ip6address** (STRING): Printable that is a loggable representation of source IPv6 address.

**ifAddrlen** (INT): Interface address length.

**lladdrlen** (INT): Address length.

## netinet6.rtr.misMth.lladrIn2

### Severity

INFORMATIONAL

### Description

This message occurs when there is a lladdrlen mismatch for the specified IPv6 address of the Router Advertising (RA) packet.

### Corrective Action

(None).

### Syslog Message

nd6\_ra\_input: lladdrlen mismatch for %s (if %d, RA packet %d).

### Parameters

**ip6address** (STRING): Printable that is a loggable representation of source IPv6 address.

**ifAddrlen** (INT): Interface address length.

**lladdrlen** (INT): Address length

## netinet6.rtr.srcAdr.invlId

### Severity

NOTICE

### Description

This message occurs during the validation of source address in the received Router Solicitation (RS) message, when it is found that the source address is a non-link local address.

### Corrective Action

(None).

### Syslog Message

nd6\_ra\_input: Source %s is not a link-local address.

### Parameters

**srcAddr** (STRING): Source IPv6 address.

# netport events

## netport.eth events

### netport.eth.discovered

**Severity**

NOTICE

**Description**

This message occurs when a physical network port has been discovered.

**Corrective Action**

(None).

**Syslog Message**

Physical network port %s (node %s) has been discovered.

**Parameters**

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

## netport.pvlan events

### netport.pvlan.created

**Severity**

NOTICE

**Description**

This message occurs when a pseudo-VLAN port has been created.

**Corrective Action**

(None).

**Syslog Message**

Pseudo-VLAN port "%s" was created on node "%s".

**Parameters**

**port** (STRING): Name of the pseudo-VLAN port.

**node** (STRING): Name of the node on which the pseudo-VLAN port resides.

# nfs events

## nfs.client events

### nfs.client.cache.modify.err

#### Severity

ERROR

#### Description

This message occurs when an NFS connected client's cache settings are changed using the REST API `"/api/protocol/nfs/connected-client-settings"`, but the value is not updated in the network component. This might lead to the cache refresh behavior not working as expected.

#### Corrective Action

Use the REST API PATCH `"/api/protocol/nfs/connected-client-settings"` to update the value.

#### Syslog Message

Connected-clients cache settings could not be updated. Error: %s.

#### Parameters

**reason** (STRING): Reason for the failure.

## nfs.krb events

### nfs.krb.lif.disabled

#### Severity

ALERT

#### Description

This message occurs when a Kerberos-enabled LIF in a Vserver shares NFS Kerberos credentials with a LIF in another Vserver. NFS Kerberos is automatically disabled by Data ONTAP® on LIFs that share Kerberos credentials across Vservers, because of the security violation this causes. Disabling Kerberos on a LIF can cause service disruption to NFS clients that use these LIFs, mounted using the `'-sec=krb5'` option.

#### Corrective Action

NFS Kerberos has been automatically disabled by Data ONTAP on the LIF because it shared Kerberos credentials with another Vserver, causing a security violation. NFS Kerberos can be reenabled using the command: `"vserver nfs kerberos interface enable"`. During the reenabling process, Kerberos credential sharing across Vservers is disallowed.

#### Syslog Message

NFS Kerberos disabled on LIF "%s" with SPN "%s" for Vserver "%s" due to a security conflict with another Vserver. NFS Kerberos can be reenabled using the command: `"vserver nfs kerberos interface enable"`.

#### Parameters

**lif** (STRING): Name of the LIF on which NFS Kerberos is disabled.  
**spn** (STRING): The service principal name associated with the LIF.  
**vserver** (STRING): Name of the Vserver that owns the LIF.

# nis events

## nis.active events

### nis.active.false.removed

**Severity**

NOTICE

**Description**

This message occurs when the last system to be upgraded in a cluster removes NIS configurations that were marked as inactive.

**Corrective Action**

(None).

**Syslog Message**

All NIS configurations with the "-active" parameter set to "false" were removed.

**Parameters**

(None).

## nis.db events

### nis.db.job.queue.failure

**Severity**

ERROR

**Description**

This message occurs when the NIS group database job queue fails.

**Corrective Action**

Use the "vserver services name-service nis-domain group-database build" command to queue the NIS database job.

**Syslog Message**

NIS group database job queue failed for Vserver %s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

## nis.group events

### nis.group.db.build.failure

**Severity**

ERROR

## Description

This message occurs when the NIS group database build fails.

## Corrective Action

Ensure that the configured NIS domain is reachable and bound. Ensure that the number of characters per record in the NIS server for group database does not exceed 1,024. Run the "vserver services name-service nis-domain group-database build" command to build the NIS group database again.

## Syslog Message

Failed to build the NIS group database for Vserver %s. Error: %s.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**error\_string** (STRING): Error encountered while building the NIS group database.

## nis.group.db.build.success

### Severity

NOTICE

### Description

This message occurs when the NIS group database build is successful.

### Corrective Action

(None).

### Syslog Message

NIS group database is successfully built for Vserver %s.

### Parameters

**vserver\_name** (STRING): Name of the Vserver.

## nis.netgrp events

### nis.netgrp.db.build.failure

#### Severity

ERROR

#### Description

This message occurs when the Network Information Service (NIS) netgroup database build fails.

#### Corrective Action

Ensure that the configured NIS domain is reachable for the SVM and bound. Ensure that the number of characters per record in the NIS server for the netgroup database does not exceed 1,024. See the ONTAP® 9 Documentation Center for managing NIS domains.

#### Syslog Message

Failed to build the NIS netgroup database for Vserver "%s". Error: %s.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**error\_string** (STRING): Error encountered while building the NIS netgroup database.

## nis.netgrp.db.build.success

### Severity

NOTICE

### Description

This message occurs when the Network Information Service (NIS) netgroup database build is successful.

### Corrective Action

(None).

### Syslog Message

NIS netgroup database is successfully built for Vserver "%s".

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

## nis.netgrp.db.job.kick.fail

### Severity

ERROR

### Description

This message occurs when the system fails to run the NIS netgroup.byhost database job. The NIS netgroup.byhost database available on the local system is not refreshed. The subsequent netgroup-by-host query might return outdated information.

### Corrective Action

Use the (privilege: advanced) "vserver services name-service nis-domain netgroup-database build" command to run the NIS netgroup database job.

### Syslog Message

NIS netgroup database job run failed for Vserver %s.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

## nis.netgrp.db.job.queue.fail

### Severity

NOTICE

### Description

This message occurs when the system fails to queue the NIS netgroup.byhost database job. The NIS netgroup.byhost database is not available on the local system. The subsequent netgroup-by-host query will be fetched from the external NIS server.

**Corrective Action**

(None).

**Syslog Message**

NIS netgroup database job queue failed for Vserver %s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

## nis.server.events

**nis.server.active****Severity**

INFORMATIONAL

**Description**

This message occurs when the Storage Virtual Machine (SVM, formerly Vserver) binds to any of the configured NIS servers.

**Corrective Action**

(None).

**Syslog Message**

SVM '%s' has successfully bound to the %s NIS server.

**Parameters**

**vserver\_name** (STRING): Name of the SVM.  
**nis\_server** (STRING): Configured NIS server.

**nis.server.inactive****Severity**

ERROR

**Description**

This message occurs when one of the NIS servers configured for a Storage Virtual Machine (SVM, formerly Vserver) cannot be contacted.

**Corrective Action**

From a UNIX ® workstation, make sure that the NIS server is responding to requests. Also make sure that the portmapper on the NIS server is responding to requests. Make sure that there are no networking issues stopping the cluster from communicating with this NIS server.

**Syslog Message**

%s NIS server configured for SVM '%s' cannot be contacted.

**Parameters**

**nis\_server** (STRING): Configured NIS server.  
**vserver\_name** (STRING): Name of the SVM.

## **nis.server.invalid.domain**

### **Severity**

ERROR

### **Description**

This message occurs when the Network Information Service (NIS) servers configured for a storage virtual machine (SVM, formerly Vserver) cannot be contacted due to "Invalid Domain" error.

### **Corrective Action**

Make sure that the NIS server with the correct domain is configured.

### **Syslog Message**

%s NIS server configured for SVM '%s' cannot be contacted due to "Invalid Domain" error.

### **Parameters**

**nis\_server** (STRING): Configured NIS server.

**vserver\_name** (STRING): Name of the SVM.

## **nis.server.not.available**

### **Severity**

ERROR

### **Description**

This message occurs when none of the NIS servers configured for a Storage Virtual Machine (SVM, formerly Vserver) can be contacted.

### **Corrective Action**

From a UNIX ® workstation, make sure that the NIS server is responding to requests. Also make sure that the portmapper on the NIS server is responding to requests. Make sure that there are no networking issues stopping the cluster from communicating with this NIS server.

### **Syslog Message**

None of the NIS servers configured for SVM '%s' can be contacted.

### **Parameters**

**vserver\_name** (STRING): Name of the SVM.



# no events

## no.halt events

### no.halt.brokenDisk

#### Severity

ALERT

#### Description

This message occurs when the option `raid.timeout` is set to 0 or if the option hasn't been set by the user, disabling automatic shutdown on a disk failure.

#### Corrective Action

Add an appropriate spare disk.

#### Syslog Message

%s in RAID group "%s" %s broken. Automatic system shutdown has been disabled.

#### Parameters

**type** (STRING): Type of disk.

**rgname** (STRING): RAID group name.

**verb** (STRING): Failed disk verb.

### no.halt.brokenDisk.notice

#### Severity

NOTICE

#### Description

This message occurs when a node is in a degraded state because of a broken disk, but the option `raid.timeout` is set to 0 or its value hasn't been modified by the user, disabling automatic shutdown on a disk failure. Shutdown is to encourage you to replace broken disks. This message occurs on a capacity-optimized platform with broken disks irrespective of the value of the option `raid.timeout`.

#### Corrective Action

Consult other logged messages or use the `"storage disk show -broken"` command to determine the identity of broken disks. Remove any broken disks. Automatic system shutdown can be enabled by altering the `"raid.timeout"` value using the `"storage raid-options modify"` command. NOTE: The option `"raid.timeout"` will not affect the automatic shutdown on a capacity-optimized platform.

#### Syslog Message

%s. automatic system shutdown is disabled. Replace the disk immediately.

#### Parameters

**type** (STRING): Type of degradation.

### no.halt.nvramLowBatt.info

**Severity**

ALERT

**Description**

This message occurs when the option `raid.timeout` has been set to 0, disabling automatic shutdown on a low NVRAM Battery condition.

**Corrective Action**

Replace the NVRAM battery. Automatic system shutdown can be enabled by altering the "`raid.timeout`" value using the "`storage raid-options modify`" command.

**Syslog Message**

Automatic system shutdown is disabled. If the NVRAM battery is dangerously low, replace it immediately.

**Parameters**

(None).

**no.halt.nvramLowBatteries****Severity**

ALERT

**Description**

This message occurs when the option `raid.timeout` is set to 0, disabling automatic shutdown on a low NVRAM Batteries condition.

**Corrective Action**

Replace the NVRAM batteries.

**Syslog Message**

NVRAM batteries are dangerously low. Automatic system shutdown is disabled. Replace the batteries immediately!

**Parameters**

(None).

**no.halt.nvramLowBattery****Severity**

ALERT

**Description**

This message occurs when the option `raid.timeout` is set to 0, disabling automatic shutdown on a low NVRAM Battery condition.

**Corrective Action**

Replace the NVRAM battery.

**Syslog Message**

NVRAM battery is dangerously low. Automatic system shutdown is disabled. Replace the battery immediately!

**Parameters**

(None).

**no.halt.nvramLowBatts.info****Severity**

NOTICE

**Description**

This message occurs when the option `raid.timeout` has been set to 0, disabling automatic shutdown on a low NVRAM Batteries condition.

**Corrective Action**

Replace the NVRAM batteries. Automatic system shutdown can be enabled by altering the "`raid.timeout`" value using the "`storage raid-options modify`" command.

**Syslog Message**

Automatic system shutdown is disabled. If the NVRAM batteries are dangerously low, we encourage you to replace it immediately.

**Parameters**

(None).

# noderun events

## noderun.login events

### noderun.login.success

#### Severity

INFORMATIONAL

#### Description

This message occurs when a cluster user successfully logs in to the node using the 'run' command.

#### Corrective Action

(None).

#### Syslog Message

'%s' logged in from cluster: '%s'.

#### Parameters

**user** (STRING): Name of the user who logged in to the node.

**cluster** (STRING): Cluster name from where the user logged in to the node.

# nodewatchdog events

## nodewatchdog.config events

### nodewatchdog.config.invalid

#### Severity

ERROR

#### Description

This message occurs when `bootarg.nodewatchdog.level` is set to an invalid value. This can only be done manually. When an invalid value is set, Node Watchdog reverts to the default behavior.

#### Corrective Action

Set `bootarg.nodewatchdog.level` to a valid value at the `LOADER>` prompt, or modify the value by using the "system node watchdog modify" command.

#### Syslog Message

The value of `bootarg.nodewatchdog.level` is set to an invalid value, "%s." Node Watchdog reverted to the default behavior (%s).

#### Parameters

**bootarg** (STRING): Value of `bootarg.nodewatchdog.level`.

**defaultVal** (STRING): Default value of `bootarg.nodewatchdog.level`.

## nodewatchdog.gfc events

### nodewatchdog.gfc.too.long

#### Severity

ALERT

#### Description

This message occurs when the storage system is low in networking buffers (mbuf) for longer than the duration threshold, which indicates abnormal circumstances. When this happens, CLI commands might hang, and clients could experience timeouts.

#### Corrective Action

If CLI commands are responding slowly or clients are timing out, reboot the node using the "reboot -node [node] -dump true" option. Also call the support team to investigate further, because this is an abnormal circumstance that might recur if not addressed. Contact NetApp technical support.

#### Syslog Message

Networking buffer depletion detected. Consumption exceeded %d%% in use and did not recede to %d%% for %d seconds.

#### Parameters

**enter\_percentage** (INT): Percentage of in-use mbufs necessary to enter the flow control state.

**exit\_percentage** (INT): Percentage of in-use mbufs necessary to exit the flow control state.

**duration** (INT): Duration for which there is a scarcity of available mbufs.

# nodewatchdog.node events

## nodewatchdog.node.failure

### Severity

EMERGENCY

### Description

This message occurs when Data ONTAP® experiences a prolonged outage of internal services critical to continued data service. The node experiencing this failure might operate in a degraded mode until the condition is addressed. Data ONTAP will attempt to recover by restarting the affected process.

### Corrective Action

The affected process may produce a core file which can be analyzed. Contact NetApp technical support if the condition persists and possible analysis of the core file.

### Syslog Message

Data ONTAP has experienced a serious internal error: %s. This might cause the node experiencing the problem to become unresponsive to data access.

### Parameters

**condition** (STRING): Condition that caused the failure.

**diagnosis** (STRING): List of system diagnoses that could cause node watchdog issues.

## nodewatchdog.node.longreboot

### Severity

ALERT

### Description

This message occurs when a node fails to reboot within the configured time allowed for rebooting.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Data ONTAP has experienced a serious internal error. The node experiencing this problem is unable to reboot within its allotted time of %d seconds causing it to be unavailable. The node has been panicked to enable it to recover.

### Parameters

**timeout** (INT): The time in seconds within which reboot did not complete.

## nodewatchdog.node.panic

### Severity

ALERT

### Description

This message occurs when Data ONTAP® experiences a prolonged outage of internal services critical to continued data service. The node has been restarted to recover from the condition.

## Corrective Action

Contact NetApp technical support for additional assistance.

## Syslog Message

Data ONTAP has experienced a serious internal error: %s. This might cause the node experiencing the problem to become unresponsive to data access. %s

## Parameters

**condition** (STRING): Condition that caused the failure.

**action** (STRING): Automatic corrective action taken (or why avoided) as a result of detecting this condition.

**diagnosis** (STRING): List of system diagnoses that could cause node watchdog issues.

## nodewatchdog.node.ucore.hung

### Severity

ALERT

### Description

This message occurs when a node fails to generate an application core within the time allotted for application coredump due to a serious internal error. The node is panicked to recover from the internal error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Unable to generate an application core for %s (pid %d) within the allotted time of %d seconds causing the application to become unavailable. The node has been panicked to recover.

### Parameters

**process\_name** (STRING): Name of the application that failed to generate core.

**process\_id** (INT): PID of the application that failed to generate core.

**timeout** (INT): Time in seconds within which the application coredump did not complete.

## nodewatchdog.resrc events

### nodewatchdog.resrc.alloc.failure

#### Severity

ALERT

#### Description

This message occurs when a critical system resource required for Data ONTAP® to continue serving data cannot be allocated for an extended period of time. Data ONTAP might operate in a degraded mode.

#### Corrective Action

If Data ONTAP operates in a degraded mode, reboot the node using the "-dump true" option. Contact NetApp technical support. When Node Watchdog is configured to "node\_restart", Data ONTAP will detect and recover from this condition automatically.

## Syslog Message

Resource allocation failures detected for %s. The percentage of %d allocation failure attempts were met or exceeded for %d seconds.

## Parameters

**resource** (STRING): Type of resource that cannot be allocated.

**percentage** (INT): Failure percentage threshold.

**duration** (INT): Duration of the resource allocation failure.

**diagnosis** (STRING): List of system diagnoses that could cause node watchdog issues.

## nodewatchdog.resrc.lng.running.crit

### Severity

ALERT

### Description

This message occurs when the count of certain critical resources internal to Data ONTAP® are consumed with long running operations. When this happens, Data ONTAP might operate in a degraded mode.

### Corrective Action

If Data ONTAP operates in a degraded mode, reboot the node using the "-dump true" option. Contact NetApp technical support. When Node Watchdog is configured to "node\_restart", Data ONTAP will detect and recover from this condition automatically.

## Syslog Message

Long running %s detected. Exceeded threshold of %d%% for %d seconds.

## Parameters

**resource** (STRING): Resource that is consumed with long running operations.

**percentage** (INT): Percentage of the long running resource.

**duration** (INT): Duration for which the resources were long running.

## nodewatchdog.svc events

### nodewatchdog.svc.rpc.noresp

### Severity

ALERT

### Description

This message occurs when a service critical to data access fails to respond to service monitoring and is restarted. Data ONTAP® might have experienced a serious error and might operate in a degraded mode.

### Corrective Action

If the message reports that the service has "restarted" then no action is required. If the status is "not restarted" or "unsuccessfully restarted," then reboot the node using the "system node reboot" command with the "-dump true" option, and then contact NetApp technical support.

## Syslog Message

The %s service internal to Data ONTAP that is required for continuing data service was unavailable. The service failed, but was %s.



## Parameters

**service** (STRING): Service that failed.

**restarted** (STRING): Whether the service was restarted. Possible values are "restarted," "not restarted," and "unsuccessfully restarted."

**diagnosis** (STRING): List of system diagnoses that could cause node watchdog issues.

# nonha events

## nonha.resvconflicthalt events

### nonha.resvConflictHalt

#### Severity

EMERGENCY

#### Description

This message occurs when the controller halts itself after it finds disk with reservation conflicts.

#### Corrective Action

Contact NetApp technical support for recovery assistance.

#### Syslog Message

A disk reservation conflict was detected on disk %s at %-20s.

#### Parameters

**disk\_name** (STRING): Name of the disk.

**time** (STRING): Time when the disk reservation conflict was detected.

**last\_update** (INT): Time passed since last rescan, in milliseconds.

# nse events

## nse.authlockout events

### nse.authLockOut

#### Severity

ERROR

#### Description

This message occurs when a NetApp® Storage Encryption (NSE) self-encrypting drive (SED) reports excessive sequential failed authentication attempts with an incorrect Authentication Key (AK) on the indicated SED authority. To prevent brute-force AK attacks, the SED will no longer accept even the correct AK. Access and control of the SED are affected. This condition applies to drives that implement the Trusted Computing Group (TCG) "Enterprise" and "Opal" specifications. Enterprise BandMaster1 authority lockout and Opal User1 authority lockout affect protection controls on user data. User data on the drive is permanently irretrievable for lockouts on these authorities. For other lockouts, some aspect of compliance with Federal Information Processing Standard (FIPS) PUB 140-2 is compromised. Enterprise "Erase Master" and Opal "Owner" authority lockouts disable cryptographic sanitize and destroy ONTAP® operations.

#### Corrective Action

The data cannot be recovered from the drive, but it might be possible to return the drive to service. To prevent lockouts, ensure that all cluster NSE nodes have all SED AKs. Display SED key IDs by using the "storage encryption disk show [-fips]" command. Verify them against those retrieved from the AK storage method: KMIP servers: Use the "security key-manager restore" command. Onboard key management: Use the "security key-manager key show" command. If the FIPS key on the drive remains at the default MSID value, for BandMaster1 or User1 authority lockout, reset the AK to its default value and simultaneously erase user data by using the "storage encryption disk sanitize" command or the (advanced privilege) "disk encrypt sanitize" command in node shell/maintenance mode. Power-cycle clears Physical Ownership lockout in some SEDs. Ensure that the SED AKs are available as above, and power-cycle the drive to see if it clears the lockout. For other lockouts, reset the SED to its as-manufactured state and erase user data by using the "storage encryption disk revert-to-original-state" command or the (advanced privilege) "disk encrypt revert\_original" command in node shell/maintenance mode with the PSID value from the device label. If the PSID is not available, replace the drive. After clearing the lockout, make the drive a spare by using the "storage disk assign" command and (advanced privilege) "storage disk unfail [-spare]" command as needed.

#### Syslog Message

Lockout on self-encrypting drive %s; security provider: %s, authority: %s, during operation "%s".

#### Parameters

**deviceName** (STRING): Name of the device.

**securityProvider** (STRING): Name of the TCG security provider process.

**authority** (STRING): Name of the TCG authority.

**operation** (STRING): Description of the operation that attempted authentication.

## nse.op events

## **nse.op.failed**

### **Severity**

ERROR

### **Description**

This message occurs when ONTAP® does not succeed in performing an operation on a NetApp® Storage Encryption (NSE) self-encrypting drive (SED). The message applies to drives that implement either the Trusted Computing Group (TCG) "Enterprise" or "Opal" specification. Some aspect of the operation of NSE drives is compromised.

### **Corrective Action**

If the error persists, replace the drive.

### **Syslog Message**

Control failure on self-encrypting drive %s; security provider: %s, authority: %s, during operation "%s".

### **Parameters**

**deviceName** (STRING): Name of the device.

**securityProvider** (STRING): Name of the TCG security provider process.

**authority** (STRING): Name of the affected TCG authority.

**operation** (STRING): Description of the operation that failed.

# nv events

## nv.data events

### nv.data.loss.possible

**Severity**

NOTICE

**Description**

This message occurs when a Cloud Volumes ONTAP node configured for high write speed unexpectedly shuts down, possibly resulting in a loss of data. This is expected because the NVRAM data does not have a chance to be flushed during the unforeseen shutdown.

**Corrective Action**

(None).

**Syslog Message**

An unexpected shutdown occurred while in high write speed mode, which possibly caused a loss of data.

**Parameters**

(None).

## nv.extension events

### nv.extension.disabled

**Severity**

INFORMATIONAL

**Description**

This message occurs when NVRAM segments of a nvlog client are disabled in a NVRAM section.

**Corrective Action**

(None).

**Syslog Message**

%s segments have been disabled in %s NVRAM section.

**Parameters**

**cid** (STRING): Nvlog client ID.

**nv\_section** (STRING): The NVRAM section in which the nvlog client segments are disabled.

### nv.extension.enabled

**Severity**

INFORMATIONAL

## Description

This message occurs when NVRAM segments of a nvlog client are enabled in a NVRAM section.

## Corrective Action

(None).

## Syslog Message

%s segments have been enabled in %s NVRAM section.

## Parameters

**cid** (STRING): Nvlog client ID.

**nv\_section** (STRING): The NVRAM section in which nvlog client segments are enabled.

## nv.extinction.failed

### Severity

ERROR

### Description

This message occurs when nvlog client segments cannot be enabled in a partner NVRAM section following a controller failover. When this event occurs, the system continues to operate, but with diminished performance due to the increased failover load coupled with decreased NVRAM capacity for logging.

### Corrective Action

For a System ID-related failure message, verify that each member node in this configuration is assigned a valid and unique system ID by using the "system show -fields systemid" command. Do not physically move a controller or NVRAM card in an HA or MetroCluster(tm) configuration. A failure message of "Configuration Recovery Needed" means that a reboot is needed to recover persistent MetroCluster configuration at a switchover site.

### Syslog Message

Client segments cannot be enabled in the %s NVRAM section, because of %s.

### Parameters

**nv\_section** (STRING): NVRAM section in which nvlog client segments cannot be enabled.

**failure\_reason** (STRING): Reason client segments cannot be enabled in this partner NVRAM section.

## nv.fake events

### nv.fake

### Severity

NOTICE

### Description

This message occurs when the nonvolatile (NV) driver detects that system memory is being used to simulate NVRAM.

### Corrective Action

Reset the 'fake-nvram?' boot environment variable from the boot loader: 'unsetenv fake-nvram?'

### **Syslog Message**

%u MB system memory being used to simulate NVRAM.

### **Parameters**

**memory\_size** (INT): Memory size (in MB) being used to simulate NVRAM.

## **nv.fio events**

### **nv.fio.write.err**

#### **Severity**

ERROR

#### **Description**

This message occurs when the write to VNVRAM fails and the retry count has been exhausted.

#### **Corrective Action**

(None).

### **Syslog Message**

The write to VNVRAM failed with error %u; retry limit exhausted.

### **Parameters**

**error** (INT): Error returned by the last write.

### **nv.fio.write.retry**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the write to VNVRAM fails. The write will be retried.

#### **Corrective Action**

(None).

### **Syslog Message**

The write to VNVRAM failed with error %u; will retry.

### **Parameters**

**error** (INT): Error returned by the write.

## **nv.flash events**

### **nv.flash.predict.fail**

#### **Severity**

ALERT

## Description

This message occurs when ONTAP® determines that failure is forthcoming or has occurred for the associated flash device on the nonvolatile memory (NVRAM) module.

## Corrective Action

Replace the device to prevent future failures.

## Syslog Message

%s

## Parameters

**reason** (STRING): Reason for failure.

## nv.flash.unable.to.monitor

### Severity

ALERT

## Description

This message occurs when the ONTAP® module that is responsible for monitoring the health of the NVRAM device cannot extract health information. This failure might indicate the NVRAM device is not operating properly.

## Corrective Action

Use the "system environment sensors show" command to verify the device's status. If the failure still exists, contact technical support for assistance.

## Syslog Message

Unable to monitor NVRAM device "%s".

## Parameters

**device** (STRING): Name of the device.

## nv.none events

### nv.none

### Severity

EMERGENCY

## Description

This message occurs when the system cannot detect a nonvolatile (NV) device.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

No NV device has been detected.



## Parameters

(None).

# nv.partner.events

## nv.partner.disabled

### Severity

INFORMATIONAL

### Description

This message occurs when checking the partner nonvolatile RAM (NVRAM) and the partner NVRAM is disabled. This event most commonly occurs during the takeover or switchover process after the partner has gone down.

### Corrective Action

(None).

### Syslog Message

NVRAM validate: %s NVRAM was disabled.

### Parameters

**host** (STRING): Name of the partner host.

## nv.partner.nvcheck.failure

### Severity

ALERT

### Description

This message occurs when nonvolatile RAM (NVRAM) inconsistencies are detected during takeover. Takeover is continued only if the option for consistent NVRAM mirroring is disabled or if the takeover is forced.

### Corrective Action

When the partner node recovers, NVRAM will be reinitialized to a consistent state. The error string provides additional information to assist NetApp technical support in identifying the inconsistency and subsystem that caused the failure.

### Syslog Message

NVRAM validate: %s

### Parameters

**error** (STRING): Error string describing which subsystem caused the failure and what the mismatching values are.

## nv.partner.size.mismatch

### Severity

ALERT

## Description

This message occurs when the HA pair nonvolatile RAM (NVRAM) sizes do not match. In this case, an HA pair cannot be formed.

## Corrective Action

Check to make sure that the local NVRAM hardware size matches the partner NVRAM hardware size. Verify that the same bootargs are in use on both nodes. One or both nodes might need to be rebooted. Contact NetApp technical support if the problem persists.

## Syslog Message

CF monitor: NVRAM size mismatch between nodes detected; node requires %llu bytes, partner exported %llu bytes.

## Parameters

**localNVSize** (LONGINT): Size of the local NVRAM.  
**partnerNVSize** (LONGINT): Size of the partner NVRAM.

## nv.partner.takeover.continued

### Severity

ALERT

### Description

This message occurs when takeover proceeds even when nonvolatile RAM (NVRAM) inconsistencies are detected in the HA pair. This most commonly occurs when the NVRAM mirror consistency has been disabled or if the takeover is forced.

### Corrective Action

(None).

### Syslog Message

NVRAM validate: Error encountered (%s), but continuing takeover; some data might be lost.

### Parameters

**error** (STRING): Error string.

## nv.rename events

### nv.rename.err

### Severity

INFORMATIONAL

### Description

This message occurs when there is an error while renaming the files specified in `/etc/nvfail_rename`.

### Corrective Action

(None).

## Syslog Message

nvfail\_rename: %s: %s.

## Parameters

**fileName** (STRING): File name.

**error** (STRING): Error string.

## nv.rename.files

### Severity

INFORMATIONAL

### Description

This message occurs while renaming the files specified in `/etc/nvfail_rename`.

### Corrective Action

(None).

## Syslog Message

Renaming files specified in `/etc/nvfail_rename`.

## Parameters

(None).

## nv.rename.longName

### Severity

INFORMATIONAL

### Description

This message occurs while renaming the files specified in `/etc/nvfail_rename`, when the file name is too long.

### Corrective Action

(None).

## Syslog Message

nvfail\_rename: %s: File name too long.

## Parameters

**fileName** (STRING): File name.

## nv.switchover events

### nv.switchover.continued

### Severity

ALERT

**Description**

This message occurs when switchover proceeds even when nonvolatile RAM (NVRAM) inconsistencies are detected in the DR partner sections.

**Corrective Action**

(None).

**Syslog Message**

NVRAM validate: Error encountered (%s), but continuing switchover; some data might be lost.

**Parameters**

**error** (STRING): Error string.

# nvd events

## nvd.biodone events

### nvd.biodone.bio.flush.error

**Severity**

NOTICE

**Description**

This message occurs when there is an error during a block I/O flush operation.

**Corrective Action**

Check the host logs for any drive errors related to the device reported in this event.

**Syslog Message**

Block I/O flush operation (bio\_flush) for opcode: 0x%x failed for device: %s.

**Parameters**

**opcode** (INHEX): SCSI opcode.

**dev\_path** (STRING): Device path.

### nvd.biodone.error.code

**Severity**

NOTICE

**Description**

This message occurs when there is an error during an I/O operation.

**Corrective Action**

Check the host logs for any drive errors related to the device reported in this event.

**Syslog Message**

I/O for opcode: 0x%x returned with error code: %d for device: %s.

**Parameters**

**opcode** (INHEX): SCSI opcode.

**bio\_error** (INT): Error code.

**dev\_path** (STRING): Device path.

# nvdimm events

## nvdimm.excessive events

### nvdimm.excessive.cecc.errs

#### Severity

ALERT

#### Description

This message occurs when excessive correctable error-correcting code (ECC) memory errors are reported by the NVDIMM controller.

#### Corrective Action

Use the "system environment sensors show -hidden true" clustershell command to view the current "NV Correctable ECC Count" value. Contact NetApp technical support for further assistance.

#### Syslog Message

Excessive correctable ECC errors in the NVDIMM in slot-%d. Count: %d.

#### Parameters

**slot** (INT): NVDIMM slot location.

**count** (INT): CECC count.

## nvdimm.flash events

### nvdimm.flash.lifetime.low

#### Severity

ALERT

#### Description

This message occurs when the NVDIMM device exceeds its warning threshold. This issue typically indicates that the NVDIMM hardware is approaching the end of its service life and should be replaced soon.

#### Corrective Action

Contact NetApp technical support for assistance with replacing the NVDIMM hardware.

#### Syslog Message

Service lifetime of NVDIMM device in slot-%d is low at %d%%.

#### Parameters

**slot** (INT): NVDIMM slot location.

**percent** (INT): Lifetime percentage.

## nvdimm.fw events

## **nvdimm.fw.upgrade.failure**

### **Severity**

ERROR

### **Description**

This message occurs when NVDIMM controller firmware failed to upgrade.

### **Corrective Action**

Power-cycle the controller and reboot ONTAP to attempt the firmware update again. If the error persists, contact NetApp technical support for assistance.

### **Syslog Message**

NVDIMM controller in slot-%d failed to upgrade firmware because of %s.

### **Parameters**

**slot** (INT): NVDIMM slot location.

**error** (STRING): Reason for upgrade failure, which could be: 1. the image file does not exist or is corrupted 2. it is unable to enter update mode 3. it is unable to clear update mode 3. the firmware image transfer failed 4. of a CRC mismatch with the transferred image 5. it failed to validate or commit the image 6. it failed to restart after the image upgrade

## **nvdimm.fw.upgrade.succeed**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVDIMM controller successfully upgrades its own firmware image..

### **Corrective Action**

(None).

### **Syslog Message**

NVDIMM controller firmware in slot-%d was upgraded to %s.

### **Parameters**

**slot** (INT): NVDIMM slot location.

**version** (STRING): New firmware version.

## **nvdimm.nvemmm events**

### **nvdimm.nvemmm.destage.succeed**

### **Severity**

NOTICE

### **Description**

This message occurs when all NVDIMM devices have copied nonvolatile data to flash memory.

**Corrective Action**

(None).

**Syslog Message**

NVMEM subsystem successfully copied nonvolatile data to flash memory on NVDIMM device in %ld seconds.

**Parameters**

**seconds** (INT): Time taken by the destage operation.

**nvdimm.nvemmm.restage.succeed****Severity**

NOTICE

**Description**

This message occurs when all NVDIMM devices have restored nonvolatile data from flash memory.

**Corrective Action**

(None).

**Syslog Message**

NVMEM subsystem successfully restored nonvolatile data from flash memory on NVDIMM device in %ld seconds.

**Parameters**

**seconds** (INT): Time taken by the restage operation.

**nvdimm.nvmem events****nvdimm.nvmem.checkfail****Severity**

EMERGENCY

**Description**

This message occurs when NVDIMM-based NVMEM checks fail.

**Corrective Action**

For all errors, check that the right type of NVDIMM is properly installed at the designated DIMM slot.

**Syslog Message**

NVMEM subsystem configuration check failed because %s. To prevent NVMEM data loss, halt the system.

**Parameters**

**error** (STRING): Description of the error, which could be: 1. device in slot-N is not accessible. 2. device in slot-N is missing. 3. device in slot-N is not NVDIMM. 4. device in slot-N is the wrong size (N); expected size (N)



## **nvdimm.nvmem.destage.failure**

### **Severity**

ALERT

### **Description**

This message occurs when NVDIMM fails to copy nonvolatile data to the flash on the NVDIMM-device.

### **Corrective Action**

This type of error indicates that the volatile data on the local NVMEM has been lost. A storage system in an HA configuration might recover the NVMEM content from its partner. To prevent future possible data loss, take the appropriate actions based on the error: 1. If the failure is caused by a malfunctioning battery, contact NetApp technical support for assistance with battery replacement. 2. If the failure is caused by an NVDIMM intermittent hardware failure, contact NetApp technical support for assistance.

### **Syslog Message**

NVMEM subsystem failed to copy nonvolatile data to flash memory on NVIDMM device because %s.

### **Parameters**

**error** (STRING): Reason for destage failure could be: (1) SDRAM is not in self-refresh mode (2) energy source (battery) is not functioning properly (3) of insufficient battery (4) destage process timeout (5) an intermitten hardware error results partial data saved (6) of a hardware error (7) of a firmware error (8) of a flash error

## **nvdimm.nvmem.initfail**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the nvdimm-based NVMEM subsystem is not initialized succesfully.

### **Corrective Action**

For all errors, check if the error is persistent or intermittent by power cycling the controller and rebooting ONTAP. Note: If the last reboot is a dirty shutdown, set the bootarg.init.battery\_override to force system to replay the NVLOG first when booting. Then, gracefully shut down and power cycle the controller.. If the failure is persistent: (a) For error (1) (2) (3) (5) (6) (7) (8), contact NetApp technical support for assistance with replacement.of NVDIMM or upgrade NVDIMM firmware. (b) For error (4), b.1 If failure cause is battery, contact NetApp technical support for assistance with replacement of the battery. b.2 For other failures, contact NetApp technical support for assistance with replacement of NVDIMM. (c) For error (9), contact NetApp technical support for assistance with replacement of the controller. (d) For error (10), verify the NVDIMM is installed properly and check if the correct BIOS is installed.

### **Syslog Message**

NVMEM subsystem initialization failed because %s. To prevent NVMEM data loss, halt the system.

### **Parameters**

**error** (STRING): Description of the error, which could be: 1.NVDIMM device in slot-N not accessible. 2.NVDIMM device in slot-N not ready. 3.NVDIMM device in slot-N is busy. 4.NVDIMM device in slot-N lost persistence due to [reason] 5.NVDIMM device in slot-N has worn-out flash (N percent remaining) 6.NVDIMM device in slot-N failed to configure energy source 7.NVDIMM device in slot-N failed to configure saveNPin. 8.NVDIMM device in slot-N failed to enable data preservation. 9.failed to enable battery.

10.NVDIMM memory address not mapped contiguously.

## **nvdimm.nvmem.restage.failure**

### **Severity**

ALERT

### **Description**

This message occurs when nonvolatile data on the NVDIMM is either corrupt or not accessible. Because these errors cannot be corrected, nonvolatile data in flash memory is lost.

### **Corrective Action**

Frequent occurrence of these errors indicates that the NVDIMM is experiencing persistent intermittent hardware failure. Contact NetApp technical support for assistance with replacement.

### **Syslog Message**

NVMEM subsystem failed to restore nonvolatile data from flash memory on the NVDIMM module because of %s.

### **Parameters**

**error** (STRING): Reason for restage failure could be: (1) SDRAM is not in self-refresh mode (2) of SDRAM uncorrectable error (3) restage process timeout (4) of a hardware error

# nvlog events

## nvlog.replay events

### nvlog.replay.source

#### Severity

NOTICE

#### Description

This message occurs on Cloud Volumes ONTAP platforms to identify the source of the nonvolatile log (nvlog) replay that follows.

#### Corrective Action

(None).

#### Syslog Message

nvlog replay data source is %s.

#### Parameters

**Source** (STRING): Source of the nvlog replay data.

# nvme events

## nvme.boot events

### nvme.boot.fw.update.complete

#### Severity

ALERT

#### Description

This message occurs when the NVMe boot device firmware has been updated, but requires a system reboot or takeover/giveback to activate the new firmware.

#### Corrective Action

Reboot or perform a takeover/giveback of the system as soon as possible.

#### Syslog Message

The NVMe boot device has completed the firmware update. The system requires a reboot or takeover/giveback to activate new firmware.

#### Parameters

(None).

## nvme.card events

### nvme.card.degraded

#### Severity

ERROR

#### Description

This message occurs when the NVMe caching module has gone into a degraded state due to a warning reported by the drive. Performance, integrity, or reliability may be compromised.

#### Corrective Action

Replace the NVMe caching module.

#### Syslog Message

NVMe caching module in slot %s with serial number %s, and model number %s is degraded with status 0x%llx.

#### Parameters

**slot\_subslot** (STRING): PCI slot of the NVMe device.

**serial** (STRING): Serial number of the controller.

**model** (STRING): Model number of the controller.

**reason** (LONGINTEX): Internal reason why the card is degraded.

## **nvme.card.removed**

### **Severity**

ERROR

### **Description**

This message occurs when an NVMe drive has been "hot-removed" from the system (while the system was powered on). The caching module is unavailable and system performance will be degraded.

### **Corrective Action**

Insert an NVMe caching module into the system.

### **Syslog Message**

NVMe caching module in slot %s with serial number %s, and model number %s was removed from the system.

### **Parameters**

**slot\_subslot** (STRING): PCI slot of the NVMe device.

**serial** (STRING): Serial number of the controller.

**model** (STRING): Model number of the controller.

## **nvme.fw events**

### **nvme.fw.activate.complete**

#### **Severity**

NOTICE

#### **Description**

This message occurs when new firmware has been successfully activated on the drive and running normally.

#### **Corrective Action**

(None).

#### **Syslog Message**

NVMe caching module in physical slot %s with serial number %s, and model number %s successfully replaced firmware version %s with %s in firmware slot %u.

#### **Parameters**

**slot\_subslot** (STRING): PCI slot of the NVMe device.

**serial** (STRING): Serial number of the controller.

**model** (STRING): Model number of the controller.

**old** (STRING): Version of firmware that has been replaced.

**new** (STRING): Version of firmware that was activated.

**fw\_slot** (INT): Firmware slot that the new version was programmed to.

**reset** (INT): An indication on whether the device had to be reset or not in order to activate the firmware.

### **nvme.fw.activate.failed**

## Severity

ERROR

## Description

This message occurs when there was an issue while activating the new firmware on the the NVMe drive. The caching module will continue normal operation with obsolete firmware.

## Corrective Action

The caching module is still completely functional, but it is running with obsolete firmware. Contact NetApp technical support for alternative solutions.

## Syslog Message

NVMe caching module in slot %s with serial number %s and model number %s failed to activate firmware. Error code: 0x%llx.

## Parameters

**slot\_subslot** (STRING): PCI slot of the NVMe device.  
**serial** (STRING): Serial number of the controller.  
**model** (STRING): Model number of the controller.  
**code** (LONGINTEX): Internal error code.

## nvme.fw.program.failed

### Severity

ERROR

### Description

This message occurs when there was an issue while downloading new firmware to the NVMe drive. The caching module will continue normal operation with obsolete firmware.

### Corrective Action

The caching module is still completely functional however it is running with obsolete firmware. Contact NetApp technical support for alternative solutions.

### Syslog Message

NVMe caching module in slot %s with serial number %s, and model number %s failed to download firmware version %s to the drive with error code 0x%llx.

### Parameters

**slot\_subslot** (STRING): PCI slot of the NVMe device.  
**serial** (STRING): Serial number of the controller.  
**model** (STRING): Model number of the controller.  
**version** (STRING): Version of firmware attempted to download.  
**code** (LONGINTEX): Internal error code.

## nvme.link events

### nvme.link.disabled.error

**Severity**

ERROR

**Description**

This message occurs when the NVMe driver disables the PCIe link due to excessive errors.

**Corrective Action**

Remove or replace the NVMe SSD.

**Syslog Message**

PCIe link disabled for NVMe SSD in slot %d due to excessive errors.

**Parameters**

**slot** (INT): Slot number of associated NVMe SSD.

**nvme.link.error****Severity**

ERROR

**Description**

This message occurs when the NVMe SSD PCIe link does not initialize.

**Corrective Action**

Remove or replace the NVMe SSD.

**Syslog Message**

PCIe link initialization error for NVMe SSD in slot %d.

**Parameters**

**slot** (INT): Slot number of associated NVMe SSD.

**nvme.ssd events****nvme.ssd.attach.error****Severity**

ERROR

**Description**

This message occurs when the NVMe drive encounters an error before the drive is reported to upper layers. The drive will not be discovered, and will not appear in sysconfig output.

**Corrective Action**

The drive could not be attached. Remove or replace the drive.

**Syslog Message**

NVMe SSD encountered error: "%s" on slot %d. Could not attach drive.

## Parameters

**error** (STRING): Type of error that occurred.

**slot** (INT): Slot number of associated NVMe SSD.



# nvmem events

## nvmem.battery events

### nvmem.battery.capacity.low

#### Severity

EMERGENCY

#### Description

This message occurs when the NVMEM battery lacks the capacity to preserve the NVMEM contents for the required minimum of 72 hours. The system is at the risk of data loss if the power fails. This message repeats every hour while the problem continues and the system shuts down in 24 hours if automatic recharging of the battery does not restore its charge.

#### Corrective Action

Correct any environmental problems, such as chassis over-temperature. The battery charges automatically. If the capacity is not restored in several hours, replace the battery pack. If the problem persists, replace the controller module.

#### Syslog Message

The NVMEM battery lacks the capacity to preserve the NVMEM contents.

#### Parameters

(None).

### nvmem.battery.capacity.low.warn

#### Severity

INFORMATIONAL

#### Description

This message occurs when the NVMEM battery capacity is below normal.

#### Corrective Action

(None).

#### Syslog Message

The NVMEM battery capacity is below normal.

#### Parameters

(None).

### nvmem.battery.capacity.normal

#### Severity

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery capacity is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery capacity is normal.

**Parameters**

(None).

**nvmem.battery.capLow****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery capacity is low. To prevent data loss, the system schedules a clean shutdown if the problem is not corrected.

**Corrective Action**

Correct any environmental problems, such as chassis overheating. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after the system is powered on, replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery capacity is low (%s).

**Parameters**

**reading** (STRING): Reading that triggered this message.

**nvmem.battery.capLowCrit****Severity**

EMERGENCY

**Description**

This message occurs when the NVMEM battery capacity is critically low. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

**Corrective Action**

Correct any environmental problems, such as overheating. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after the system is powered on, replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery capacity is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

## Parameters

**reading** (STRING): Reading that triggered this message.  
**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.capLowWarn**

### Severity

ERROR

### Description

This message occurs when the NVMEM battery capacity is below normal.

### Corrective Action

Correct any environmental problems, such as overheating. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after the system is powered on, replace the NVMEM battery within three months. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery capacity is below normal (%s).

## Parameters

**reading** (STRING): Reading that triggered this message.

## **nvmem.battery.capNormal**

### Severity

NOTICE

### Description

This message occurs when the NVMEM battery capacity is normal.

### Corrective Action

(None).

### Syslog Message

The NVMEM battery capacity is normal.

## Parameters

(None).

## **nvmem.battery.current.high**

### Severity

EMERGENCY

### Description

This message occurs when the NVMEM battery current is excessively high and the system will shut down.

### Corrective Action

First, correct any environmental problems, such as chassis over-temperature. If the NVMEM battery current is still too high, replace the battery pack. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery current is excessively high and the system will shut down.

**Parameters**

(None).

**nvmem.battery.current.high.warn****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery current is above normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery current is above normal.

**Parameters**

(None).

**nvmem.battery.current.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery current is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery current is normal.

**Parameters**

(None).

**nvmem.battery.currentHigh****Severity**

EMERGENCY

**Description**

This message occurs when the NVMEM battery charging current is critically high. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Replace the NVMEM battery. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery current is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

### Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## nvmem.battery.currentHiWarn

### Severity

ALERT

### Description

This message occurs when the NVMEM battery charging current is above normal. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Replace the NVMEM battery. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery current is above normal (%s). To prevent data loss, the system will shut down %s.

### Parameters

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## nvmem.battery.currentLow

### Severity

EMERGENCY

### Description

This message occurs when the NVMEM battery detects a short circuit. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Replace the controller module.

### Syslog Message

The NVMEM battery has a short circuit (%s). To prevent data loss, the system will shut down in %d minutes.

### Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.currentLowWarn**

### **Severity**

ALERT

### **Description**

This message occurs when the NVMEM battery is discharging when it should not. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

Replace the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery current is below normal (%s). To prevent data loss, the system will shut down %s.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvmem.battery.discFET.normal**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMEM battery discharge field-effect transistor (FET) is on.

### **Corrective Action**

(None).

### **Syslog Message**

The NVMEM battery discharge FET is on.

### **Parameters**

(None).

## **nvmem.battery.discFET.off**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVMEM battery discharge field-effect transistor (FET) is off. To prevent data loss, the system shuts down.

### **Corrective Action**

Shut down the system and power-cycle the controller. If the discharge FET remains off, replace the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery discharge FET is off. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvmem.battery.end\_of\_life.high****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery-cycle count indicates that the battery has reached its anticipated life expectancy.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery has reached its anticipated life expectancy.

**Parameters**

(None).

**nvmem.battery.end\_of\_life.high.warn****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery-cycle count indicates that the battery is nearing its anticipated life expectancy.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery is nearing its anticipated life expectancy.

**Parameters**

(None).

**nvmem.battery.end\_of\_life.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery-cycle count indicates that the battery is well below its anticipated life expectancy.

**Corrective Action**

(None).

## Syslog Message

The NVMEM battery is well below its anticipated life expectancy.

## Parameters

(None).

## **nvmem.battery.fccLow**

### Severity

ALERT

### Description

This message occurs when the NVMEM battery full-charge capacity is low. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Use the "system battery show" SP CLI command to get the battery manufactured date. If the battery is newer than three years old, correct any chronic environmental problems, such as overheating. If the full-charge capacity is still below normal, contact technical support for assistance with replacement of the NVMEM battery. If the problem persists, replace the controller module.

## Syslog Message

The NVMEM battery full-charge capacity is low (%s). To prevent data loss, the system will shut down %s.

## Parameters

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvmem.battery.fccLowCrit**

### Severity

EMERGENCY

### Description

This message occurs when the NVMEM battery full-charge capacity is critically low. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Use the "system battery show" SP CLI command to get the battery manufactured date. If the battery is newer than three years old, correct any chronic environmental problems, such as overheating. If the full-charge capacity is still below normal, contact technical support for assistance with replacement of the NVMEM battery. If the problem persists, replace the controller module.

## Syslog Message

The NVMEM battery full-charge capacity is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

## Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.



## **nvmem.battery.fccLowWarn**

### **Severity**

ALERT

### **Description**

This message occurs when the NVMEM battery full-charge capacity is below normal.

### **Corrective Action**

Use the "system battery show" SP CLI command to get the battery manufactured date. If the battery is newer than three years old, correct any chronic environmental problems, such as overheating. If the full-charge capacity is still below normal, contact technical support for assistance with replacement of the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery full-charge capacity is below normal (%s).

### **Parameters**

**reading** (STRING): Reading that triggered this message.

## **nvmem.battery.fccNormal**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMEM battery full-charge capacity is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVMEM battery full-charge capacity is normal.

### **Parameters**

(None).

## **nvmem.battery.FET.normal**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMEM battery charge field-effect transistor (FET) is on.

### **Corrective Action**

(None).

### **Syslog Message**

The NVMEM battery charge FET is on.

**Parameters**

(None).

**nvmem.battery.FET.off****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery charge field-effect transistor (FET) is off. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

**Corrective Action**

Shut down the system and power-cycle the controller. If the charge FET remains off, replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery charge FET is off. To prevent data loss, the system will shut down %s.

**Parameters**

**countdown** (STRING): Remaining time before chassis shutdown.

**nvmem.battery.noCharge****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery requests a charge but does not receive it. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

**Corrective Action**

Pull the controller from the slot and ensure that the NVMEM battery connector is seated properly. If the problem occurs again, replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery is requesting to be charged but the charger is not charging the battery. To prevent data loss, the system will shut down %s.

**Parameters**

**countdown** (STRING): Remaining time before chassis shutdown.

**nvmem.battery.normalCharge****Severity**

NOTICE

**Description**

This message occurs when the NVMEM battery charging status is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery charging status is normal.

**Parameters**

(None).

**nvmem.battery.notPresent****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery is not present. To prevent data loss, the system shuts down.

**Corrective Action**

While the system is shut down, pull the controller from the slot and ensure that the NVMEM battery is present. If the NVMEM battery is already present, ensure that the connector is seated properly. Push the controller back in and power on the system. If this does not correct the problem, replace the NVMEM battery. If the system still does not detect the battery, replace the controller module.

**Syslog Message**

The NVMEM battery is not present. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvmem.battery.packInvalid****Severity**

EMERGENCY

**Description**

This message occurs when an unsupported NVMEM battery pack is installed in the system. To prevent data loss, the system initiates a clean shutdown if the problem is not corrected.

**Corrective Action**

Shut down the system and replace the NVMEM battery with a correct battery pack.

**Syslog Message**

The incorrect NVMEM battery is installed. Expected part number is %s. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**partnum** (STRING): Valid battery part number.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.packInvalidHrs**

### **Severity**

EMERGENCY

### **Description**

This message occurs when an unsupported NVMEM battery pack is installed in the system. To prevent data loss, the system initiates a clean shutdown if the problem is not corrected.

### **Corrective Action**

Shut down the system and replace the NVMEM battery with a correct battery pack.

### **Syslog Message**

The incorrect NVMEM battery is installed. Expected part number is %s. To prevent data loss, the system will shut down in %s.

### **Parameters**

**partnum** (STRING): Valid battery part number.

**remaining\_time** (STRING): Number of hours until the system shuts down.

## **nvmem.battery.packValid**

### **Severity**

NOTICE

### **Description**

This message occurs when a valid NVMEM battery pack is present and readable.

### **Corrective Action**

(None).

### **Syslog Message**

A valid NVMEM battery pack is present.

### **Parameters**

(None).

## **nvmem.battery.powerFault**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVMEM battery is not receiving power. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

Pull the controller from the slot and ensure that the NVMEM battery connector is seated properly. Push the controller back in and boot the system. If the NVMEM battery is still not receiving power, replace the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery is not receiving power. To prevent data loss, the system will shut down in %d minutes.

### **Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.powerNormal**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMEM battery power is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVMEM battery power is normal.

### **Parameters**

(None).

## **nvmem.battery.sensor.unread**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the battery state of the battery-backed memory (NVMEM) is unknown. One of the battery sensors is not readable.

### **Corrective Action**

Shut down the system and power-cycle the controller module. If the problem is not corrected, replace the battery. If the sensor is still not readable, replace the controller module.

### **Syslog Message**

The battery state of the battery-backed memory (NVMEM) %s is not readable.

### **Parameters**

**sensor\_name** (STRING): Sensor name.

## **nvmem.battery.temp.high**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVMEM battery is too hot and the system is at a high risk of data loss if power fails. The system will shut down.

**Corrective Action**

If the system is excessively warm, allow it to cool gradually. If the NVMEM battery temperature reading is still too high, replace the battery pack. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery is too hot and the system is at a high risk of data loss if power fails. The system will shut down to prevent data loss.

**Parameters**

(None).

**nvmem.battery.temp.high.warn****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery temperature high.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery temperature is high.

**Parameters**

(None).

**nvmem.battery.temp.low****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery is too cold and the system is at a high risk of data loss if power fails.

**Corrective Action**

If the system is excessively cold, allow it to warm gradually. If the NVMEM battery temperature reading is still too low, replace the battery pack. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery is too cold and the system is at a high risk of data loss if power fails.

**Parameters**

(None).

**nvmem.battery.temp.low.warn**

**Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery temperature is low.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery temperature is low.

**Parameters**

(None).

**nvmem.battery.temp.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery temperature is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery temperature is normal.

**Parameters**

(None).

**nvmem.battery.tempHigh****Severity**

EMERGENCY

**Description**

This message occurs when the NVMEM battery temperature is critically high. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

**Corrective Action**

Correct any environmental problems, such as a overheating. If the NVMEM battery temperature is still above normal, replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery temperature is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.tempLow**

### **Severity**

ALERT

### **Description**

This message occurs when the NVMEM battery is too cold. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

If the system is excessively cold, allow it to warm gradually. If the NVMEM battery temperature is still too low, replace the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery is too cold (%s). To prevent data loss, the system will shut down %s.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvmem.battery.unread**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVMEM battery state cannot be determined because the sensor is unreadable. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

Pull the controller from the slot and ensure that the NVMEM battery connector is seated properly. If the problem is not corrected, replace the NVMEM battery. If the sensor is still not readable, replace the controller module.

### **Syslog Message**

The NVMEM battery state cannot be determined because the '%s' sensor is unreadable. To prevent data loss, the system will shut down in %d minutes.

### **Parameters**

**sensor\_name** (STRING): Sensor name.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvmem.battery.unreadHrs**

### **Severity**

EMERGENCY



## Description

This message occurs when the NVMEM battery state cannot be determined because the sensor is unreadable. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

## Corrective Action

Pull the controller from the slot and ensure that the NVMEM battery connector is seated properly. If the problem is not corrected, replace the NVMEM battery. If the sensor is still not readable, replace the controller module.

## Syslog Message

The NVMEM battery state cannot be determined because the '%s' sensor is unreadable. To prevent data loss, the system will shut down in %s.

## Parameters

**sensor\_name** (STRING): Sensor name.

**remaining\_time** (STRING): Remaining time before chassis shutdown.

## **nvmem.battery.voltage.high**

### Severity

EMERGENCY

### Description

This message occurs when the NVMEM battery voltage is excessively high and the system will shut down.

### Corrective Action

First, correct any environmental problems, such as chassis over-temperature. If the NVMEM battery voltage is still too high, replace the battery pack. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery voltage is excessively high and the system will shut down.

### Parameters

(None).

## **nvmem.battery.voltage.high.warn**

### Severity

INFORMATIONAL

### Description

This message occurs when the NVMEM battery voltage is above normal.

### Corrective Action

(None).

### Syslog Message

The NVMEM battery voltage is above normal.

**Parameters**

(None).

**nvmem.battery.voltage.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM battery voltage is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM battery voltage is normal.

**Parameters**

(None).

**nvmem.battery.voltageHigh****Severity**

EMERGENCY

**Description**

This message occurs when the NVMEM battery voltage is critically high. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

**Corrective Action**

Replace the NVMEM battery. If the problem persists, replace the controller module.

**Syslog Message**

The NVMEM battery voltage is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvmem.battery.voltageHiWarn****Severity**

ALERT

**Description**

This message occurs when the NVMEM battery voltage is above normal. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Correct any environmental problems, such as overheating. If the NVMEM battery voltage is still above normal, replace the NVMEM battery. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery voltage is above normal (%s). To prevent data loss, the system will shut down %s.

### Parameters

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## nvmem.battery.voltageLow

### Severity

EMERGENCY

### Description

This message occurs when the NVMEM battery voltage is critically low. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### Corrective Action

Correct any environmental problems, such as overheating. If the NVMEM battery voltage is still critically low, replace the NVMEM battery. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery voltage is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

### Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## nvmem.battery.voltageLowWarn

### Severity

ALERT

### Description

This message occurs when the NVMEM battery voltage is below normal. To prevent data loss, the system schedules a clean shutdown if the problem is not corrected.

### Corrective Action

Correct any environmental problems, such as overheating. If the NVMEM battery voltage is still below normal 30 minutes after environmental problems have been corrected, replace the NVMEM battery. If the problem persists, replace the controller module.

### Syslog Message

The NVMEM battery voltage is below normal (%s).

### Parameters

**reading** (STRING): Reading that triggered this message.

## **nvmem.battery.wrongCharge**

### **Severity**

ALERT

### **Description**

This message occurs when the NVMEM battery charger charges the battery without a request to do so. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

Replace the NVMEM battery. If the problem persists, replace the controller module.

### **Syslog Message**

The NVMEM battery charger is charging the battery even though the battery is not requesting to be charged. To prevent data loss, the system will shut down %s.

### **Parameters**

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvmem.voltage events**

### **nvmem.voltage.high**

#### **Severity**

EMERGENCY

#### **Description**

This message occurs when the NVMEM supply voltage is high and the system is at a high risk of data loss if power fails. The system will shut down.

#### **Corrective Action**

Correct any environmental or battery problems. If the problem persists, replace the controller module.

#### **Syslog Message**

The NVMEM supply voltage is high and the system is at a high risk of data loss if power fails. The system will shut down to prevent data loss.

#### **Parameters**

(None).

### **nvmem.voltage.high.warn**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the NVMEM supply voltage is above normal.

#### **Corrective Action**

(None).

**Syslog Message**

The NVMEM supply voltage is above normal.

**Parameters**

(None).

**nvmem.voltage.low****Severity**

ALERT

**Description**

This message occurs when the NVMEM supply voltage is low and the system is at a high risk of data loss if power fails.

**Corrective Action**

First, correct any environmental or battery problems. If the problem continues, replace the controller module.

**Syslog Message**

The NVMEM supply voltage is low and the system is at a high risk of data loss if power fails.

**Parameters**

(None).

**nvmem.voltage.low.warn****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM supply voltage is below normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM supply voltage is below normal.

**Parameters**

(None).

**nvmem.voltage.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVMEM supply voltage is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVMEM supply voltage is normal.

**Parameters**

(None).

# nvmens events

## nvmens.512blksz events

### NVMeNS.512blksz.repl.unsupp

#### Severity

ERROR

#### Description

This message occurs when replication of a volume containing NVMe namespaces that use 512-byte block sizes fails because the destination does not support NVMe namespaces with this block size.

#### Corrective Action

Replicate volumes containing NVMe namespaces with 512-byte block sizes only to destinations that support namespaces that use this block size.

#### Syslog Message

Logical replication failed for volume %s (DSID %llu): No 512-byte block size NVMe namespace support on the destination.

#### Parameters

**volume\_name** (STRING): Name of the source volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the source volume.

## nvmens.destroy events

### NVMeNS.destroy

#### Severity

NOTICE

#### Description

This message occurs when an NVMe namespace has been destroyed.

#### Corrective Action

(None).

#### Syslog Message

NVMe namespace %s, vol %s (DSID %llu) was destroyed (UUID: %s).

#### Parameters

**NVMeNS\_path** (STRING): Path for the NVMe namespace that was destroyed.  
**volume\_name** (STRING): Name of the containing volume.  
**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.  
**object\_uuid** (STRING): UUID of the NVMe namespace.  
**object\_type** (STRING): Type of object under notification.

# **nvmens.offline events**

## **NVMeNS.offline**

### **Severity**

NOTICE

### **Description**

This message occurs when an NVMe namespace has been brought offline manually.

### **Corrective Action**

(None).

### **Syslog Message**

NVMe namespace %s, vol %s (DSID %llu) was brought offline (UUID: %s).

### **Parameters**

**path** (STRING): Path for the NVMe namespace that has been brought offline.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the NVMe namespace.

**object\_type** (STRING): Type of object under notification.

# **nvmens.online events**

## **NVMeNS.online**

### **Severity**

NOTICE

### **Description**

This message occurs when an NVMe namespace is brought online manually.

### **Corrective Action**

(None).

### **Syslog Message**

NVMe namespace %s, vol %s (DSID %llu) was brought online (UUID: %s).

### **Parameters**

**path** (STRING): Path for the NVMe namespace that was brought online.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the NVMe namespace.

**object\_type** (STRING): Type of object under notification.

# **nvmens.out events**



## NVMeNS.out.of.space

### Severity

EMERGENCY

### Description

This message occurs when an NVMe namespace has been brought offline because of a write failure due to lack of space.

### Corrective Action

Add space to the volume, and then bring the NVMe namespace online using the "vserver nvme namespace modify" command.

### Syslog Message

NVMe namespace %s, vol %s (DSID %llu) is out of space (UUID: %s).

### Parameters

**path** (STRING): Path for the NVMe namespace that has been brought offline.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the containing volume.

**object\_uuid** (STRING): UUID of the NVMe namespace.

**object\_type** (STRING): Type of object under notification.

## nvmens.repl events

### NVMeNS.repl.unsupp.ver

### Severity

ERROR

### Description

This message occurs when replication of a volume containing NVMe namespaces has failed because the destination does not support NVMe namespaces.

### Corrective Action

Replicate volumes with NVMe namespaces only to destinations that support NVMe namespaces.

### Syslog Message

Logical replication failed for volume %s (DSID %llu): No NVMe namespace support on the destination.

### Parameters

**volume\_name** (STRING): Name of the source volume.

**volume\_dsid** (LONGINT): Data set ID (DSID) of the source volume.

## nvmens.smbc events

### NVMeNS.SMBC.vol.unsupp

### Severity

ERROR

**Description**

This message occurs when an NVMe namespace is partially restored to a SnapMirror Business Continuity (SMBC) volume because the SMBC volume does not support NVMe namespace.

**Corrective Action**

Restore of NVMe namespace to the SMBC volume is not supported although it is partially created on the SMBC volume as part of restore. The NVMe namespace needs to be manually deleted. Contact NetApp technical support for assistance.

**Syslog Message**

NVMe namespace "%s" is not supported on SMBC volume "%s" (DSID %llu).

**Parameters**

**path** (STRING): Path for the NVMe namespace that has been restored.

**volume\_name** (STRING): Name of the containing volume.

**volume\_dsids** (LONGINT): Dataset ID (DSID) of the containing volume.

# nvmeof events

## nvmeof.cabling events

### nvmeof.cabling.clear

**Severity**

NOTICE

**Description**

This message occurs when all of the cabling issues to the shelf have been corrected.

**Corrective Action**

(None).

**Syslog Message**

The storage cabling issues on node "%s" have been corrected.

**Parameters**

**node\_name** (STRING): Node name connected to the cable.

### nvmeof.cabling.error

**Severity**

ERROR

**Description**

This message occurs when there are errors with the cabling to the shelf.

**Corrective Action**

Run the "storage check-cables" node shell command to list the errors. Follow the cabling guide to correct the issues.

**Syslog Message**

There is a storage cabling issue on node "%s". Run the "storage check-cables" node shell command to list the errors.

**Parameters**

**node\_name** (STRING): Node name connected to the cable.

## nvmeof.hostport events

### nvmeof.hostport.clear

**Severity**

NOTICE

**Description**

This message occurs when a node shelf is no longer cabled incorrectly from a port that is not a storage

port. Indicates condition has cleared.

### **Corrective Action**

(None).

### **Syslog Message**

Shelf on node %s with SN %s is no longer cabled incorrectly from a port that is not a storage port.

### **Parameters**

**node\_name** (STRING): Host name connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.hostport.error**

### **Severity**

ERROR

### **Description**

This message occurs when the host port is incorrectly cabled to the shelf.

### **Corrective Action**

Reconnect the host port to a Storage Ethernet port.

### **Syslog Message**

Port %s is not configured as a Storage Ethernet port on node %s. It is incorrectly cabled to Shelf with SN %s.

### **Parameters**

**host\_port** (STRING): Host port connected to the cable.

**node\_name** (STRING): Host name connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.mixedports events**

### **nvmeof.mixedports.clear**

### **Severity**

NOTICE

### **Description**

This message occurs when there is no longer a nonrecommended mix of a dedicated port and a nondedicated port cabled from a node to a shelf.

### **Corrective Action**

(None).

### **Syslog Message**

There is no longer a nonrecommended mix of a dedicated port and a nondedicated port cabled from node %s to shelf with SN %s.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.mixedports.error**

### Severity

ERROR

### Description

This message occurs when both primary and secondary host ports are incorrectly cabled to a shelf. This cabling might compromise physical data link redundancy and inhibit system performance.

### Corrective Action

Connect either primary dedicated ports or secondary nondedicated ports to the shelf, but not both port types.

### Syslog Message

There is a nonrecommended mix of a dedicated port and a nondedicated port cabled from node %s to Shelf with SN %s.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.onlyoneport events**

### **nvmeof.onlyoneport.clear**

### Severity

NOTICE

### Description

This message occurs when there is no longer only one port cabled from a node to a shelf.

### Corrective Action

(None).

### Syslog Message

There is no longer only one port cabled from node %s to shelf with SN %s.

## Parameters

**node\_name** (STRING): Host name connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

### **nvmeof.onlyoneport.error**

### Severity

ERROR

## Description

This message occurs when host port is incorrectly cabled to a shelf.

## Corrective Action

The host port is incorrectly cabled to shelf. two ports must be connected to ensure redundancy.

## Syslog Message

There is only one port cabled from node %s to Shelf with SN %s.

## Parameters

**node\_name** (STRING): Host name connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

# nvmeof.port events

## nvmeof.port.modify

### Severity

NOTICE

### Description

This message occurs when the mode of an NVMe-oF port is changed.

### Corrective Action

(None).

### Syslog Message

Changing NVMe-oF port %s to %s mode.

### Parameters

**port** (STRING): Name of port being configured.

**mode** (STRING): Mode in which to configure port: network or storage.

# nvmeof.portioma events

## nvmeof.portioma.clear

### Severity

NOTICE

### Description

This message occurs when there are no longer two ports cabled incorrectly from a node to top module (A) on a shelf.

### Corrective Action

(None).

### Syslog Message

Two ports are no longer cabled incorrectly from node %s to top module (A) on shelf with SN %s.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.portioma.error**

### Severity

ERROR

### Description

This message occurs when there are two ports cabled to the top module (A) on a shelf.

### Corrective Action

Connect two ports to different modules to provide redundancy.

### Syslog Message

There are two ports cabled from node %s to top module (A) on Shelf with SN %s. One of these cables should be removed.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

## **nvmeof.portiomb events**

### **nvmeof.portiomb.clear**

### Severity

NOTICE

### Description

This message occurs when there are no longer two ports cabled incorrectly from a node to bottom module (B) on a shelf.

### Corrective Action

(None).

### Syslog Message

Two ports are no longer cabled incorrectly from node %s to bottom module (B) on shelf with SN %s.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

### **nvmeof.portiomb.error**

### Severity

ERROR

## Description

This message occurs when there are two ports cabled to the bottom module (B) on a shelf.

## Corrective Action

Connect two ports to different modules to provide redundancy.

## Syslog Message

There are two ports cabled from node %s to bottom module (B) on Shelf with SN %s. One of these cables should be removed.

## Parameters

**node\_name** (STRING): Host port connected to the cable.

**shelf\_sn** (STRING): Serial number of the shelf connected to the cable.

# nvmeof.rdma.events

## nvmeof.rdma.disconnect.event

### Deprecated

Deprecated as of version 9.13.1.

### Severity

NOTICE

### Description

This message occurs when the NVMe-oF driver detects an RDMA disconnect. There might be a decrease in performance while reconnecting.

### Corrective Action

(None).

### Syslog Message

NVMe-oF RDMA disconnect on subnqn %s, controller ID %d, qpair ID %d.

### Parameters

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

## nvmeof.rdma.disconnect.status

### Deprecated

Deprecated as of version 9.13.1.

### Severity

NOTICE

### Description

This message occurs when the NVMeoF driver detects a command that finishes with a disconnect status. There might be a decrease in performance while reconnecting.



### Corrective Action

(None).

### Syslog Message

NVMe-oF disconnect on subnqn %s, controller ID %d, qpair ID %d.

### Parameters

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

## nvmeof.subsystem events

### nvmeof.subsystem.add

#### Severity

NOTICE

#### Description

This message occurs when an NVMe-oF subsystem is added.

#### Corrective Action

(None).

#### Syslog Message

NVMe-oF subsystem added at address %s.

#### Parameters

**address** (STRING): IP address of subsystem.

### nvmeof.subsystem.delete

#### Severity

NOTICE

#### Description

This message occurs when an NVMe-oF subsystem is deleted.

#### Corrective Action

(None).

#### Syslog Message

NVMe-oF subsystem deleted at address %s.

#### Parameters

**address** (STRING): IP address of subsystem.

## nvmeof.timeout events

## **nvmeof.timeout**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMeoF driver detects a timeout on a request. There might be a decrease in performance while retrying.

### **Corrective Action**

(None).

### **Syslog Message**

Timeout on subnqn %s, controller ID %d, qpair ID %d, sequence number %d.

### **Parameters**

**subnqn** (STRING): Subsystem NVMe qualified name.

**ctrlr\_id** (INT): Controller ID.

**qpair** (INT): Qpair ID.

**cid** (INT): NVMe-oF CID.

# nvmf events

## nvmf.graceperiod events

### nvmf.graceperiod.active

#### Severity

ALERT

#### Description

This message occurs daily while the NVMe over Fabrics (NVMe-oF) protocol is in use and the license grace period is active. NVMe-oF functionality requires a license in this version of ONTAP® software after the license grace period expires.

#### Corrective Action

Contact your sales representative to obtain an NVMe-oF license and add it to the cluster, or remove all NVMe-oF configuration from the cluster.

#### Syslog Message

The NVMe-oF feature requires a license in this version of ONTAP. NVMe-oF functionality will be disabled in %d days (%s) unless a license is added to the cluster.

#### Parameters

**days\_remaining** (INT): Days until the license grace period expires.

**expiration\_date** (STRING): Date the license grace period expires.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type is CLUSTER.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID is that of the cluster.

### nvmf.graceperiod.expired

#### Severity

ALERT

#### Description

This message occurs when the NVMe over Fabrics (NVMe-oF) license grace period expires and NVMe-oF functionality is disabled. NVMe-oF functionality requires a license in this version of ONTAP® software after the license grace period expires.

#### Corrective Action

Contact your sales representative to obtain an NVMe-oF license, and then add it to the cluster.

#### Syslog Message

The NVMe-oF feature requires a license in this version of ONTAP and the grace period has expired. NVMe-oF functionality will be disabled until a license is added to the cluster.

#### Parameters

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type is CLUSTER.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID is that of the cluster.

## **nvmf.graceperiod.start**

### **Severity**

ALERT

### **Description**

This message occurs when an NVMe over Fabrics (NVMe-oF) configuration is detected during upgrade to ONTAP® 9.5 software. NVMe-oF functionality requires a license in this version of ONTAP after the license grace period expires.

### **Corrective Action**

Contact your sales representative to obtain an NVMe-oF license, and then add it to the cluster.

### **Syslog Message**

The NVMe-oF feature requires a license in this version of ONTAP. NVMe-oF functionality will be disabled in %d days (%s) unless a license is added to the cluster.

### **Parameters**

**days\_remaining** (INT): Days until the license grace period expires.

**expiration\_date** (STRING): Date the license grace period expires.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type is CLUSTER.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID is that of the cluster.

## **nvmf.mgmt events**

### **nvmf.mgmt.op.failed**

#### **Severity**

EMERGENCY

#### **Description**

This message occurs when a management operation fails in an NVMe-oF subsystem. The NVMe-oF subsystem on this node will be out of sync with the subsystems on the other nodes in the cluster. This can result in the failure of administrative commands and IO.

#### **Corrective Action**

To recover from this condition in HA pair configurations, take over the node with the problem by using the "storage failover takeover" command, and then give back the node by using the "storage failover giveback" command. To recover from this condition in non-HA configurations, reboot the nodes having this condition by using the "system node reboot" command.

#### **Syslog Message**

NVMe-oF operation "%s" failed (reason=%s).

#### **Parameters**

**operation** (STRING): Management operation that failed.

**reason** (STRING): Reason for the failure.

## **nvmf.mgmt.port.cmd.full**

### **Severity**

ERROR

### **Description**

This message occurs when the NVMe-oF command fan-in on a physical port exceeds the port's capabilities. At this point, NVMe-oF I/O throughput and latency could severely degrade.

### **Corrective Action**

To recover from this condition, evaluate the host fan-in on the port, and then take one of the following actions: - Reduce the number of hosts that are logging in to this port. - Reduce the number of NVMe-oF subsystems that are being accessed by the hosts. - Reduce the queue depths and number of queues for hosts connecting to this port. After completing the appropriate mitigation step, you can see the number of command slots currently being used via the fcp\_port CM object. Example: statistics show -object fcp\_port -counter nvmf\_command\_slots -instance port.portname -raw

### **Syslog Message**

NVMe-oF port name "%s" command slots full (current number provisioned = %d).

### **Parameters**

**portname** (STRING): Port name on which the command slots are full.

**command\_slots** (INT): Current number of command slots provisioned on the port.

## **nvmf.mgmt.port.cmd.threshold**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVMe-oF command fan-in on a physical port starts approaching levels close to the port's capabilities.

### **Corrective Action**

Consider adding more NVMe-oF ports or access the NVMe-oF subsystems via other ports instead of using this port. You can see the number of command slots currently being used via the fcp\_port CM object. Example: statistics show -object fcp\_port -counter nvmf\_command\_slots -instance port.portname -raw

### **Syslog Message**

NVMe-oF port name "%s" command slots high threshold reached (current number provisioned = %d).

### **Parameters**

**portname** (STRING): Port name on which the command slots are being filled up.

**command\_slots** (INT): Current number of command slots provisioned on the port.

## **nvmf.mgmt.roce.port.cq.full**

### **Severity**

ERROR

## Description

This message occurs when the NVMe/RoCE completion queue entry utilization exceeds the port's capabilities. At this point, NVMe/RoCE I/O throughput and latency could severely degrade.

## Corrective Action

To recover from this condition, evaluate the host fan-in on the port, and then take one of the following actions: - Reduce the number of hosts that are logging in to this port. - Reduce the number of NVMe-oF subsystems that are being accessed by the hosts. - Reduce the queue depths and number of queues for hosts connecting to this port.

## Syslog Message

Completion queue entries on NVMe/RoCE port "%s" are exhausted (current number used is %d).

## Parameters

**portname** (STRING): Port name on which the completion queue entries are exhausted.

**cq\_entries** (INT): Current number of completion queue entries used on the port.

# nvmf.qos events

## nvmf.qos.mismatched.policy

### Severity

ERROR

### Description

This message occurs when an NVMe-oF target subsystem contains namespaces that have different QoS policies.

### Corrective Action

Determine which volumes exist within the named subsystem NQN by using the "vserver nvme subsystem map show" command. Verify that these volumes share the same QoS policy by using the "volume show -fields name,qos-policy-group" command.

### Syslog Message

NVMe-oF target subsystem NQN %s contains %d namespaces that do not have the same QoS policy.

### Parameters

**subsystem\_nqn** (STRING): NVMe Qualified Name (NQN) of the subsystem.

**ns\_count** (INT): Number of namespaces that do not match.

## nvmf.qos.viodet.regfail

### Severity

ERROR

### Description

This message occurs when the Quality-of-Service (QoS) subsystem fails to register for updates during initialization due to an internal error (e.g. there is not enough system memory available). As a result, workloads might not meet their minimum throughput or their service-level objectives (SLOs).

### Corrective Action

The registration will continue to be retried. If this message continues to occur, perform a planned takeover/giveback sequence to reset the reporting node nondisruptively by using the "takeover -ofnode (reporting\_node)" and "giveback -ofnode (reporting\_node)" commands. If the problem persists, contact Contact NetApp technical support..

### Syslog Message

Failed to register for violation detection updates. QoS min-throughput might operate in a limited capacity.

### Parameters

**errorCode** (INT): Error code that was returned during registration.

## nvmf.remote events

### nvmf.remote.status.cb

#### Severity

NOTICE

#### Description

This message occurs when NVMe over Fabric(NVMeOF) receives a status callback from the connection manager for the Cluster Session Manager(CSM) session is used for remote IO.

#### Corrective Action

(None).

#### Syslog Message

NVMeOF remote status callback endpoint=%s, status=%s.

#### Parameters

**endpoint** (STRING): CSM session endpoint on which the status callback is received.

**status** (STRING): Reason for the status callback.

## nvmf.roce events

### nvmf.roce.lif.create.failed

#### Severity

ERROR

#### Description

This message occurs when there was an internal error while creating a NVMe over RoCE logical network interface (LIF).

#### Corrective Action

To recover from this condition, use the "network interface modify" command to change the status-admin of the LIF to "down" and then back to "up" again.

## Syslog Message

NVMe/RoCE LIF "%s" create operation failed because of an internal error.

## Parameters

**lifname** (STRING): Name of the NVMe/RoCE LIF that encountered the failure.

# nvmf.tcp events

## nvmf.tcp.lif.create.fail

### Severity

ERROR

### Description

This message occurs when there is an internal error while creating a NVMe-over-TCP (NVMe/TCP) logical network interface (LIF).

### Corrective Action

To recover from this condition, use the "network interface modify" command to change the state of "status-admin" of the LIF to "down" and then back to "up" again.

## Syslog Message

NVMe/TCP LIF "%s" create operation failed because of the following internal error: "%d".

## Parameters

**lifname** (STRING): Name of the NVMe/TCP LIF that encountered the failure.

**errorCode** (INT): Error code returned during the creation of the LIF.

## nvmf.tcp.trans.init.fail

### Severity

ERROR

### Description

This message occurs when there is an internal error during the NVMe-over-TCP (NVMe/TCP) transport initialization.

### Corrective Action

NVMe-oF logs at "/mroot/etc/log/rastrace" need to be collected for further analysis. Trigger an ALL AutoSupport® message by using the following command: "autosupport invoke -node \* -message "nvmf.tcp.transport.init.failed" -type all". Also trigger the diagnostic AutoSupport function by using the following command "autosupport invoke-diagnostic -node \* -message "nvmf.tcp.transport.init.failed" -subsystems nvmf". Then contact NetApp technical support support.

## Syslog Message

NVMe/TCP transport initialization failed because of the following internal error: "%d".

## Parameters

**errorCode** (INT): Error code returned during transport initialization.



# nvpm events

## nvpm.extend events

### nvpm.extend.failed

#### Severity

ERROR

#### Description

This message occurs when nvlog client segments cannot be enabled in a partner NVRAM section following a controller failover or MetroCluster® switchover. When this event occurs, the system continues to operate, but with diminished performance due to the increased failover load coupled with decreased NVRAM capacity for logging.

#### Corrective Action

Perform a storage system giveback operation or a metrocluster switchback operation to restore operating performance.

#### Syslog Message

Client segments cannot be enabled due to a %s.

#### Parameters

**event** (STRING): Event that triggered the extension.

**current\_state** (INT): (debug only) Current Non-volatile RAM Partition Manager (NVPM) state.

**expected\_state** (INT): (debug only) Expected NVPM state.

**reason** (INT): (debug only) Reason for the extension failure.

# nvrām events

## nvrām.battery events

### nvrām.battery.capacity.low.critical

#### Severity

EMERGENCY

#### Description

This message occurs when the NVRAM battery capacity is critically low. To prevent data loss, the system will shut down.

#### Corrective Action

Correct any environmental problems, such as chassis over-temperature. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after system power is on, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

#### Syslog Message

The NVRAM battery capacity is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

#### Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

### nvrām.battery.capacity.low.warn

#### Severity

NOTICE

#### Description

This message occurs when the NVRAM battery capacity is below normal.

#### Corrective Action

Correct any environmental problems, such as chassis over-temperature. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after system power is on, replace the NVRAM battery within three months. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

#### Syslog Message

The NVRAM battery capacity is below normal (%s).

#### Parameters

**reading** (STRING): Reading that triggered this message.

### nvrām.battery.capacity.normal

**Severity**

INFORMATIONAL

**Description**

This message occurs when the NVRAM battery capacity is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVRAM battery capacity is normal.

**Parameters**

(None).

**nvr.am.battery.capacityLow****Severity**

ERROR

**Description**

This message occurs when the NVRAM battery capacity is low. To prevent data loss, the system will schedule a clean chassis shutdown if the problem is not corrected.

**Corrective Action**

Correct any environmental problems, such as chassis over-temperature. The battery charges automatically. If the capacity is not restored automatically within 30 minutes after system power is on, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery capacity is low (%s).

**Parameters**

**reading** (STRING): Reading that triggered this message.

**nvr.am.battery.chargeFET.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVRAM battery charge FET is on.

**Corrective Action**

(None).

**Syslog Message**

The NVRAM battery charge FET is on.

**Parameters**

(None).

**nvr.am.battery.chargeFET.off****Severity**

ERROR

**Description**

This message occurs when the NVRAM battery charge FET is off. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

**Corrective Action**

Shut down the system and power-cycle the controller. If the charge FET remains off, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery charge FET is off. To prevent data loss, the system will shut down %s.

**Parameters**

**countdown** (STRING): Remaining time before chassis shutdown.

**nvr.am.battery.charging.nocharge****Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery is requesting to be charged but the charger is not charging the battery. To prevent data loss, the system will shut down.

**Corrective Action**

Reseat the NVRAM battery connector. If the problem occurs again, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery is requesting to be charged but the charger is not charging the battery. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvr.am.battery.charging.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVRAM battery charging status is normal.

**Corrective Action**

(None).

**Syslog Message**

The NVRAM battery charging status is normal.

**Parameters**

(None).

**nvr.am.battery.charging.wrongcharge****Severity**

ERROR

**Description**

This message occurs when the NVRAM battery charger is charging the battery even though the battery is not requesting to be charged. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

**Corrective Action**

Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery charger is charging the battery even though the battery is not requesting to be charged. To prevent data loss, the system will shut down %s.

**Parameters**

**countdown** (STRING): Remaining time before chassis shutdown.

**nvr.am.battery.current.high****Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery charging current is critically high. To prevent data loss, the system will shut down.

**Corrective Action**

Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery current is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.current.high.warn**

### **Severity**

ERROR

### **Description**

This message occurs when the NVRAM battery charging current is above normal. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### **Corrective Action**

Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery current is above normal (%s). To prevent data loss, the system will shut down %s.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvram.battery.current.low**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVRAM battery detects a short circuit. To prevent data loss, the system will shut down.

### **Corrective Action**

Contact NetApp technical support for assistance with identification and replacement of the appropriate component.

### **Syslog Message**

The NVRAM battery has a short circuit (%s). To prevent data loss, the system will shut down in %d minutes.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.current.low.warn**

### **Severity**

ERROR

### **Description**

This message occurs when the NVRAM battery is discharging when it should not. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### **Corrective Action**

Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery current is below normal (%s). To prevent data loss, the system will shut down %s.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvr.am.battery.current.normal**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery current is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery current is normal.

### **Parameters**

(None).

## **nvr.am.battery.dischargeFET.normal**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery discharge FET is on.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery discharge FET is on.

### **Parameters**

(None).

## **nvr.am.battery.dischargeFET.off**

### **Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery discharge FET is off. To prevent data loss, the system will shut down.

**Corrective Action**

Shut down the system and power-cycle the controller. If the discharge FET remains off, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery discharge FET is off. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvr.am.battery.endOfLife.high****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVRAM battery-cycle count indicates that the battery has reached its anticipated life expectancy.

**Corrective Action**

Replace the NVRAM battery in the next three months.

**Syslog Message**

The NVRAM battery has reached its anticipated life expectancy.

**Parameters**

(None).

**nvr.am.battery.endOfLife.normal****Severity**

INFORMATIONAL

**Description**

This message occurs when the NVRAM battery-cycle count indicates that the battery is well below its anticipated life expectancy.

**Corrective Action**

(None).

**Syslog Message**

The NVRAM battery is well below its anticipated life expectancy.

**Parameters**

(None).



## **nvram.battery.fault**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVRAM battery is reporting a fatal fault condition. To prevent data loss, the system will shut down.

### **Corrective Action**

Correct any environmental problems, such as chassis over-temperature. If the battery still reports a fatal fault, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery is reporting a fatal fault condition. To prevent data loss, the system will shut down in %d minutes.

### **Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.fault.cleared**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery internal fault condition has been cleared.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery internal fault condition has been cleared.

### **Parameters**

(None).

## **nvram.battery.fault.warn**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVRAM battery is reporting a non-fatal fault condition.

### **Corrective Action**

Correct any environmental problems, such as chassis over- temperature. If the condition persists, replace the NVRAM battery in the next three months.

## Syslog Message

The NVRAM battery is reporting a non-fatal fault condition.

## Parameters

(None).

## **nvr.am.battery.fcc.low**

### Severity

ERROR

### Description

This message occurs when the NVRAM battery full-charge capacity is low. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### Corrective Action

Use the "sysconfig -M" command in 7-Mode or the "run local sysconfig -M" command in clustered Data ONTAP® to get the battery manufacturer date. If the battery is newer than three years old, correct any chronic environmental problems, such as chassis over-temperature. Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

## Syslog Message

The NVRAM battery full-charge capacity is low (%s). To prevent data loss, the system will shut down %s.

## Parameters

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvr.am.battery.fcc.low.critical**

### Deprecated

Deprecated as of version 9.2.0 to address the EMS ID length violation.

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM battery full-charge capacity is critically low. To prevent data loss, the system will shut down.

### Corrective Action

Use the "sysconfig -M" command in 7-Mode or the "run local sysconfig -M" command in clustered Data ONTAP® to get the battery manufacturer date. If the battery is newer than three years old, correct any chronic environmental problems, such as chassis over-temperature. Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

## Syslog Message

The NVRAM battery full-charge capacity is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

## Parameters

**reading** (STRING): Reading that triggered this message.  
**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.fcc.low.warn**

### Severity

NOTICE

### Description

This message occurs when the NVRAM battery full-charge capacity is below normal.

### Corrective Action

Replace the NVRAM battery in the next three months. Use the "sysconfig -M" command in 7-Mode or the "run local sysconfig -M" command in clustered Data ONTAP® to get the battery manufacturer date. If the battery is newer than three years old, correct any chronic environmental problems, such as chassis over-temperature.

### Syslog Message

The NVRAM battery full-charge capacity is below normal (%s).

## Parameters

**reading** (STRING): Reading that triggered this message.

## **nvram.battery.fcc.normal**

### Severity

INFORMATIONAL

### Description

This message occurs when the NVRAM battery full-charge capacity is normal.

### Corrective Action

(None).

### Syslog Message

The NVRAM battery full-charge capacity is normal.

## Parameters

(None).

## **nvram.battery.fccLowCrit**

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM battery full-charge capacity is critically low. To prevent data loss, the system will shut down.

### Corrective Action

Use the "system node run -node local sysconfig -M" command to obtain the battery manufacturing date. If the battery is over three years old, replacement might be necessary. Contact NetApp technical support for troubleshooting assistance.

### Syslog Message

The NVRAM battery full-charge capacity is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

### Parameters

**reading** (STRING): Reading that triggered this message.  
**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.fcmtto.warn**

### Severity

NOTICE

### Description

This message occurs when the NVRAM battery fast-charge timeout expires. The battery will not be charged until this condition is corrected.

### Corrective Action

Shut down the controller and power cycle it from the Service Processor (SP) command line interface (CLI) with the "system power cycle" command. The fast-charge timeout will be reset after the power cycle.

### Syslog Message

The NVRAM battery has reached a charging timeout fault. The battery charger will not turn on again until the controller has been power cycled. Low capacity warning messages begin in %s weeks.

### Parameters

**num\_weeks** (STRING): Number of weeks before low capacity warning messages begin.

## **nvram.battery.not.present**

### Severity

EMERGENCY

### Description

This message occurs when the presence of the NVRAM battery cannot be detected. To prevent data loss, the system will shut down.

### Corrective Action

While the system is shut down, ensure that the NVRAM battery is present. If the NVRAM battery is already present, ensure that the connector is seated properly. Power on the system. If this does not correct the problem, replace the NVRAM battery. If the system still does not detect the battery, contact NetApp technical support.

### Syslog Message

The presence of NVRAM battery cannot be detected. To prevent data loss, the system will shut down in %d minutes.

## Parameters

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.pack.invalid**

### Severity

ERROR

### Description

This message occurs when an unsupported NVRAM battery pack is installed in the system. To prevent data loss, the system initiates a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### Corrective Action

Shut down the system and replace the NVRAM battery with a correct battery pack.

### Syslog Message

The incorrect NVRAM battery is installed. Expected part number is %s. To prevent data loss, the system will shut down in %d minutes.

## Parameters

**partnum** (STRING): Valid battery part number.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.pack.valid**

### Severity

INFORMATIONAL

### Description

This message occurs when a valid NVRAM battery pack is now present and readable.

### Corrective Action

(None).

### Syslog Message

A valid NVRAM battery pack is present.

## Parameters

(None).

## **nvram.battery.packInvalidHrs**

### Severity

EMERGENCY

### Description

This message occurs when an unsupported NVRAM battery pack is installed in the system. To prevent data loss, the system initiates a clean shutdown if the problem is not corrected.

### **Corrective Action**

Shut down the system and replace the NVRAM battery with the correct battery pack.

### **Syslog Message**

The incorrect NVRAM battery is installed. Expected part number is "%s". To prevent data loss, the system shuts down %s.

### **Parameters**

**partnum** (STRING): Valid battery part number.

**remaining\_time** (STRING): Number of hours until the system shuts down.

## **nvr.am.battery.power.fault**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVRAM battery is not receiving power. To prevent data loss, the system will shut down.

### **Corrective Action**

Ensure that the NVRAM battery connector is seated properly. Boot up the system. If the NVRAM battery is still not receiving power, replace the NVRAM battery. If the problem persists, contact NetApp technical support.

### **Syslog Message**

The NVRAM battery is not receiving power. To prevent data loss, the system will shut down in %d minutes.

### **Parameters**

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvr.am.battery.power.normal**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery power is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery power is normal.

### **Parameters**

(None).

## **nvr.am.battery.sensor.unreadable**

**Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery state cannot be determined because the sensor cannot be read. To prevent data loss, the system will shut down.

**Corrective Action**

Ensure that the NVRAM battery connector is seated properly. Power on the system. If the problem is not corrected, replace the NVRAM battery. If the sensor still cannot be read, contact NetApp technical support.

**Syslog Message**

The NVRAM battery state cannot be determined because the "%s" sensor cannot be read. To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**sensor\_name** (STRING): Sensor name.

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvram.battery.temp.high****Severity**

EMERGENCY

**Description**

This message occurs when the NVRAM battery temperature is critically high. To prevent data loss, the system will shut down.

**Corrective Action**

Correct any environmental issues, such as a high ambient temperature. If the NVRAM battery temperature is still above normal, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

**Syslog Message**

The NVRAM battery temperature is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

**Parameters**

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

**nvram.battery.temp.high.warn****Severity**

NOTICE

**Description**

This message occurs when the NVRAM battery temperature is above normal. Battery life will be severely compromised.

### **Corrective Action**

Correct any environmental issues, such as a high ambient temperature. If the NVRAM battery temperature is still above normal, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery temperature is above normal (%s).

### **Parameters**

**reading** (STRING): Reading that triggered this message.

## **nvr.am.battery.temp.low**

### **Severity**

ERROR

### **Description**

This message occurs when the NVRAM battery is too cold. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### **Corrective Action**

If the system is excessively cold, allow it to warm gradually. If the NVRAM battery temperature is still too low, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery is too cold (%s). To prevent data loss, the system will shut down %s.

### **Parameters**

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvr.am.battery.temp.low.warn**

### **Severity**

NOTICE

### **Description**

This message occurs when the NVRAM battery temperature is below normal.

### **Corrective Action**

Allow the system to warm gradually. If the NVRAM battery temperature is still low, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery temperature is below normal (%s).

### **Parameters**

**reading** (STRING): Reading that triggered this message.



## **nvram.battery.temp.normal**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery temperature is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery temperature is normal.

### **Parameters**

(None).

## **nvram.battery.unreadHrs**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVRAM battery state cannot be determined because the sensor is unreadable. If this issue is not resolved within the specified time, the system shuts down to prevent data loss.

### **Corrective Action**

Ensure that the NVRAM battery connector is seated properly. If the problem is not corrected, replace the NVRAM battery. If the sensor is still not readable, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery state cannot be determined because the "%s" sensor is unreadable. To prevent data loss, the system shuts down %s.

### **Parameters**

**sensor\_name** (STRING): Sensor name.

**remaining\_time** (STRING): Remaining time before chassis shutdown.

## **nvram.battery.voltage.high**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the NVRAM battery voltage is critically high. To prevent data loss, the system will shut down.

### **Corrective Action**

Replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

## Syslog Message

The NVRAM battery voltage is critically high (%s). To prevent data loss, the system will shut down in %d minutes.

## Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvr.am.battery.voltage.high.warn**

### Severity

ERROR

### Description

This message occurs when the NVRAM battery voltage is above normal. To prevent data loss, the system will initiate a clean chassis shutdown if the problem is not corrected within 24 hours of first detection.

### Corrective Action

Correct any environmental problems, such as chassis over-temperature. If the NVRAM battery voltage is still above normal, replace the NVRAM battery. If the problem returns or persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

## Syslog Message

The NVRAM battery voltage is above normal (%s). To prevent data loss, the system will shut down %s.

## Parameters

**reading** (STRING): Reading that triggered this message.

**countdown** (STRING): Remaining time before chassis shutdown.

## **nvr.am.battery.voltage.low**

### Severity

EMERGENCY

### Description

This message occurs when the NVRAM battery voltage is critically low. To prevent data loss, the system will shut down.

### Corrective Action

Correct any environmental problems, such as chassis over-temperature. If the NVRAM battery voltage is still critically low, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

## Syslog Message

The NVRAM battery voltage is critically low (%s). To prevent data loss, the system will shut down in %d minutes.

## Parameters

**reading** (STRING): Reading that triggered this message.

**num\_minutes** (INT): Number of minutes until the system shuts down.

## **nvram.battery.voltage.normal**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the NVRAM battery voltage is normal.

### **Corrective Action**

(None).

### **Syslog Message**

The NVRAM battery voltage is normal.

### **Parameters**

(None).

## **nvram.battery.voltageLowWarn**

### **Severity**

ERROR

### **Description**

This message occurs when the NVRAM battery voltage is below normal. To prevent data loss, the system will schedule a clean chassis shutdown if the problem is not corrected.

### **Corrective Action**

Correct any environmental problems, such as chassis over-temperature. If the NVRAM battery voltage is still below normal 30 minutes after environmental problems have been corrected, replace the NVRAM battery. If the problem persists, replace the FRU (field-replaceable unit) NVRAM card or PCM, where the battery is located.

### **Syslog Message**

The NVRAM battery voltage is below normal (%s).

### **Parameters**

**reading** (STRING): Reading that triggered this message.

## **nvram.excessive events**

### **nvram.excessive.cecc.errs**

### **Severity**

ALERT

### **Description**

This message occurs when excessive correctable error-correcting code (ECC) memory errors are reported by the NVRAM complex.

**Corrective Action**

Use the "system environment sensors show -hidden true" cluster shell command to view the current "NV Correctable ECC count" and "NV Correctable ECC rate" values. Contact NetApp technical support for further assistance.

**Syslog Message**

Excessive correctable ECC errors in the NVRAM card.

**Parameters**

(None).

**nvr.am.excessive.sbit.ecc.errs****Deprecated**

Deprecated as of version 9.2.0 to address the EMS ID length violation.

**Severity**

ALERT

**Description**

This event is issued to report excessive singlebit ECC errors.

**Corrective Action**

(None).

**Syslog Message**

Excessive singlebit ECC errors in the NVRAM card

**Parameters**

(None).

**nvr.am.flash.events****nvr.am.flash.spare.low****Severity**

ALERT

**Description**

This message occurs periodically when the capacity of the NVRAM flash memory spare blocks is low, before it becomes critically low. This issue typically indicates that the NVRAM hardware is approaching the end of its service life and should be replaced soon.

**Corrective Action**

Contact NetApp technical support for replacement assistance.

**Syslog Message**

NVRAM flash memory spare capacity is low: %s.

## Parameters

**problem** (STRING): Information about the NVRAM flash memory spare blocks whose capacity is low and about its condition.

# nvram.hw.events

## nvram.hw.degraded

### Severity

ERROR

### Description

This message occurs when the NVRAM hardware is degraded.

### Corrective Action

Contact NetApp technical support for replacement assistance.

### Syslog Message

NVRAM hardware is degraded: %s.

### Parameters

**reasonText** (STRING): Reason the NVRAM hardware is degraded.

## nvram.hw.initFail

### Severity

EMERGENCY

### Description

This message occurs when the Data ONTAP® NVRAM hardware fails to initialize.

### Corrective Action

Typically, this type of error is unexpected and indicates that the NVRAM hardware is failing and should be replaced. Contact NetApp technical support for assistance with the replacement.

### Syslog Message

NVRAM hardware initialization failed: %s.

### Parameters

**error** (STRING): Description of the error.

## nvram.hw.normal

### Severity

NOTICE

### Description

This message occurs when the NVRAM hardware is operating normally.

**Corrective Action**

(None).

**Syslog Message**

NVRAM hardware is operating normally.

**Parameters**

(None).

# nwd events

## nwd.dataip events

### nwd.dataip.removed

#### Severity

NOTICE

#### Description

This message occurs when a data IP address is being removed from a node. Data access through this IP address is now unavailable.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s is being removed from this node.

#### Parameters

**dataip** (STRING): The IP address that is being removed.

# oauth2 events

## oauth2.httpd events

### oauth2.httpd.restart.failed

#### Severity

ERROR

#### Description

This message occurs when the httpd process fails to restart when the OAuth 2.0 status changes.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

When the OAuth 2.0 status changed, the httpd process failed to restart for the following reason: %s.

#### Parameters

**reason** (STRING): Reason why the httpd process failed to restart when there was a change in the OAuth 2.0 status of the cluster.

## oauth2.localroles events

### oauth2.localroles.unauth

#### Severity

ERROR

#### Description

This message occurs when an authentication fails when the "use-local-roles-if-present" setting is set to "true", but the user information does not exist in ONTAP, or the scope information within the access token does not grant sufficient access in ONTAP.

#### Corrective Action

Add a local user entry using the "security login create" command, or include the relevant scope information in the access token.

#### Syslog Message

Authentication of requested user "%s" in SVM "%s" failed because the "use-local-roles-if-present" setting is set to "true"; however, the user information does not exist in ONTAP, or the scope information within the access token does not grant sufficient access in ONTAP.

#### Parameters

**user** (STRING): User for whom authentication failed.

**vserver** (STRING): SVM to which the user belongs.



## oauth2.status events

### oauth2.status.disabled

**Severity**

NOTICE

**Description**

This message occurs when the OAuth 2.0 status of the cluster is changed to DISABLED.

**Corrective Action**

(None).

**Syslog Message**

Status of OAuth 2.0 has been changed to DISABLED.

**Parameters**

(None).

### oauth2.status.enabled

**Severity**

NOTICE

**Description**

This message occurs when the OAuth 2.0 status of the cluster is changed to ENABLED.

**Corrective Action**

(None).

**Syslog Message**

Status of OAuth 2.0 has been changed to ENABLED.

**Parameters**

(None).

## oauth2.wsdb events

### oauth2.wsdb.set.failed

**Severity**

ERROR

**Description**

This message occurs when setting the SSL trusted server CA path in web services database fails. Potential reason could be that the application is busy or unresponsive. This will impact OAuth2.0 user authentication.

**Corrective Action**

Ensure all nodes in the cluster are online and healthy by running the "cluster show" command and verifying that the "Health" and "Eligibility" fields displayed as true. If true, retry the operation. If the issue persists,

contact NetApp technical support.

**Syslog Message**

Failed to set the SSL trusted server CA path in web services database.

**Parameters**

(None).

# object events

## object.store events

### object.store.available

#### Severity

NOTICE

#### Description

This message occurs when the node can connect to the object store again after an "object.store.unavailable" error.

#### Corrective Action

(None).

#### Syslog Message

Able to connect to the object store "%s" from node %s.

#### Parameters

**configname** (STRING): Object store configuration name.  
**node\_uuid** (STRING): UUID of the node.

### object.store.c2s.clock.ahead

#### Severity

ERROR

#### Description

This message occurs when the cluster has a Commercial Cloud Services (C2S) object store configured, and the clock on a node is skewed forward more than the maximum allowed number of minutes.

#### Corrective Action

Correct the clock on a node by using the "date" command. Make sure that the nodes sync their clocks through the NTP service by using the "cluster time-service ntp" commands. If the condition persists, contact NetApp technical support.

#### Syslog Message

The clock on node "%s" is ahead by more than the maximum of %d minutes. Current skew is %d:%d minutes.

#### Parameters

**node\_uuid** (STRING): UUID of the node.  
**max\_allowed\_skew\_minutes** (INT): Maximum allowed number of minutes by which the clock on a node can be skewed forward.  
**clock\_skew\_minutes** (INT): Number of minutes by which the clock on a node is skewed forward.  
**clock\_skew\_seconds** (INT): Number of seconds by which the clock on a node is skewed forward.

## object.store.c2s.clock.late

### Severity

ERROR

### Description

This message occurs when the cluster has a Commercial Cloud Services (C2S) object store configured, and the clock on a node is skewed backwards more than the maximum allowed number of minutes. Data in the C2S object store might become inaccessible for some time because the node might not update expired temporary credentials on time.

### Corrective Action

Correct the clock on a node by using the "date" command. Make sure that the nodes synchronize their clocks through the NTP service by using the "cluster time-service ntp" commands. If the condition persists, contact NetApp technical support.

### Syslog Message

The clock on node "%s" is behind by more than the maximum of %d minutes. Current skew is %d:%d minutes.

### Parameters

**node\_uuid** (STRING): UUID of the node.

**max\_allowed\_skew\_minutes** (INT): Maximum allowed number of minutes by which the clock on a node can be skewed backwards.

**clock\_skew\_minutes** (INT): Number of minutes by which the clock on a node is skewed backwards.

**clock\_skew\_seconds** (INT): Number of seconds by which the clock on a node is skewed backwards.

## object.store.c2s.notupdated

### Severity

EMERGENCY

### Description

This message occurs when the node fails to update temporary credentials required for accessing a Commercial Cloud Services (C2S) object store. Data in the C2S object store might not be accessible.

### Corrective Action

Verify that LIFs in the appropriate IPspace are online and functional by using the "network interface show" command. Also, check network connectivity to the CAP through a LIF in the appropriate IPspace by using the "network ping" command. Additionally, ensure that security certificates for the server certificate authorities (CA) and for the client are installed and valid. If the condition persists, contact NetApp technical support.

### Syslog Message

Failed to update temporary credentials that expired on %s in object store "%s" from node "%s". %s. CAP URL: %s

### Parameters

**credentials\_expiration** (STRING): Date and time when temporary credentials expired.

**object\_store\_name** (STRING): Object store configuration name.

**node\_uuid** (STRING): UUID of the node.

**reason** (STRING): Reason the temporary credentials were not updated.

**cap\_url** (STRING): Full URL of the request to a C2S Authentication Portal (CAP).

## **object.store.c2s.updated**

### **Severity**

NOTICE

### **Description**

This message occurs when the node updates temporary credentials required for accessing a Commercial Cloud Services (C2S) object store after an "object.store.c2s.notupdated" error.

### **Corrective Action**

(None).

### **Syslog Message**

Successfully updated temporary credentials for object store "%s" from node "%s". New expiration: %s, CAP URL: %s

### **Parameters**

**object\_store\_name** (STRING): Object store configuration name.

**node\_uuid** (STRING): UUID of the node.

**credentials\_expiration** (STRING): Updated date and time when temporary credentials expire.

**cap\_url** (STRING): Full URL of the request to a C2S Authentication Portal (CAP).

## **object.store.certificate.off**

### **Severity**

NOTICE

### **Description**

This message occurs when SSL/TLS certificate validation is disabled for an object store configuration.

### **Corrective Action**

Enable certificate validation by using the "storage aggregate object-store config modify -object-store-name <name> -is-certificate-validation-enabled true" command.

### **Syslog Message**

SSL/TLS certificate validation is disabled for an object store configuration %s.

### **Parameters**

**object\_store\_name** (STRING): Name of the object store configuration.

## **object.store.certificate.on**

### **Severity**

NOTICE

### **Description**

This message occurs when SSL/TLS certificate validation is enabled for an object store configuration for which it was previously disabled.

### Corrective Action

(None).

### Syslog Message

SSL/TLS certificate validation is enabled for object store configuration %s.

### Parameters

**object\_store\_name** (STRING): Name of the object store configuration.

## object.store.degraded.mirror.swapped

### Severity

NOTICE

### Description

This message occurs when a FabricPool's primary object store and its mirror object store are swapped while the mirror object store is in a degraded state.

### Corrective Action

(None).

### Syslog Message

The primary object store (%s) and the degraded mirror object store (%s) are swapped for FabricPool %s.

### Parameters

**primary\_object\_store** (STRING): Name of the primary object store before swap.

**mirror\_object\_store** (STRING): Name of the mirror object store before swap.

**aggr** (STRING): Name of the FabricPool.

## object.store.force.tier.mcc

### Severity

NOTICE

### Description

This message occurs when the "-force-tiering-on-metrocluster" parameter is changed for an object store attached to a FabricPool that is in a MetroCluster(tm) configuration.

### Corrective Action

(None).

### Syslog Message

The "-force-tiering-on-metrocluster" parameter is set to "%s" on the object store "%s" attached to FabricPool "%s".

### Parameters

**is\_force\_tier\_on\_mcc** (STRING): Current setting of the "-force-tiering-on-metrocluster" parameter (true or false).

**object\_store** (STRING): Name of the object store.

**aggr** (STRING): Name of the FabricPool.

## **object.store.full**

### **Severity**

ALERT

### **Description**

This message occurs when the node failed to write data to the specified object store because it is out of space. You may see other EMS notifications from dependent subsystems.

### **Corrective Action**

Examine the object store, check the space settings, and if necessary, create space by adding capacity to the object store.

### **Syslog Message**

Failed to write to object store %s because it is out of space.

### **Parameters**

**config\_name** (STRING): Object store configuration name.

## **object.store.mirror.added**

### **Severity**

NOTICE

### **Description**

This message occurs when the second object store is attached to the FabricPool. After mirror operation, FabricPool contain primary and mirror object store.

### **Corrective Action**

(None).

### **Syslog Message**

A second object store (%s) is added to FabricPool %s with primary object store (%s) as a mirror.

### **Parameters**

**mirror\_object\_store** (STRING): Name of the mirror object store.

**aggr** (STRING): Name of the FabricPool.

**primary\_object\_store** (STRING): Name of the primary object store.

## **object.store.mirror.swapped**

### **Severity**

NOTICE

### **Description**

This message occurs when a FabricPool's primary object store and its mirror object store are swapped.

### **Corrective Action**

(None).

## Syslog Message

The primary object store (%s) and the mirror object store (%s) are swapped for FabricPool %s.

## Parameters

**primary\_object\_store** (STRING): Name of the primary object store before swap.

**mirror\_object\_store** (STRING): Name of the mirror object store before swap.

**aggr** (STRING): Name of the FabricPool.

## object.store.mirror.unmirror

### Severity

NOTICE

### Description

This message occurs when the mirror object store is removed from a FabricPool. FabricPool is said to be unmirrored.

### Corrective Action

(None).

## Syslog Message

The mirror object store (%s) is removed from FabricPool %s with primary object store (%s).

## Parameters

**mirror\_object\_store** (STRING): Name of the mirror object store.

**aggr** (STRING): Name of the FabricPool.

**primary\_object\_store** (STRING): Name of the primary object store.

## object.store.no.proxy

### Severity

ALERT

### Description

This message occurs when one or more object store configurations have use-http-proxy set to true but no HTTP proxy is configured for the IPspace. This can lead to client I/O failures.

### Corrective Action

Create an HTTP proxy configuration for the IPspace or change use-http-proxy to false for the object store configurations associated with that IPspace.

## Syslog Message

One or more object store configurations have use-http-proxy set to true but no HTTP proxy configuration was found for the IPspace.

## Parameters

**ipspace\_name** (STRING): Name of the IPspace that does not have an HTTP proxy server configured.

## object.store.unavailable



**Severity**

EMERGENCY

**Description**

This message occurs when the node cannot connect to the object store.

**Corrective Action**

Verify that your intercluster LIF is online and functional by using the "network interface show" command. Also, check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. Additionally, ensure that the configuration of your object store has not changed, and that login and connectivity information is still accurate. If the condition persists, contact NetApp technical support.

**Syslog Message**

Unable to connect to the object store "%s" from node %s. Reason: %s.

**Parameters**

**configname** (STRING): Object store configuration name.

**node\_uuid** (STRING): UUID of the node.

**reason** (STRING): Reason for the failure to connect.

# objectattach events

## objectattach.hw events

### objectAttach.hw.unsupported

#### Severity

NOTICE

#### Description

This message occurs when an object store attach failed due to an unsupported hardware configuration.

#### Corrective Action

(None).

#### Syslog Message

Object store attach failed because node %s with model %s is not supported by FabricPools.

#### Parameters

**node** (STRING): Node that has the unsupported hardware version.

**model** (STRING): Model information.

# objstore events

## objstore.host events

### objstore.host.unresolvable

#### Severity

ERROR

#### Description

This message occurs when the object-store server host name cannot be resolved to an IP address. The object-store client will not be able to communicate with the object-store server without resolving to an IP address.

#### Corrective Action

Check the DNS configuration to verify that the host name is configured correctly with an IP address.

#### Syslog Message

Object-store server host name "%s" cannot be resolved to an IP address on node %s.

#### Parameters

**hostname** (STRING): Host name of the object-store server.

**nodeUuid** (STRING): UUID of the node.

## objstore.interclusterlifdown events

### objstore.interclusterlifDown

#### Severity

ERROR

#### Description

This message occurs when the object-store client cannot find an operational LIF to communicate with the object-store server. This node will not allow object-store client traffic until the intercluster LIF is operational.

#### Corrective Action

Use the "network interface show -role intercluster" command to check intercluster LIF status. Verify that the intercluster LIF is configured correctly and is operational. If an intercluster LIF is not configured, use the "network interface create -role intercluster" command to create it.

#### Syslog Message

Object-store client could not find an operational intercluster LIF (IPspace ID: %d) on node %s.

#### Parameters

**ipspaceID** (INT): IPspace ID of the intercluster LIF.

**nodeUuid** (STRING): UUID of the node.

# old events

## old.replica events

### old.replica.volume

#### Severity

ALERT

#### Description

This message occurs when a wafiron operation determines that a volume served as a SnapMirror® destination for a previous release of Data ONTAP®.

#### Corrective Action

Retry the wafiron operation after the SnapMirror relationship is updated for this volume.

#### Syslog Message

Wafiron cannot start on aggregate %s because aggregate contains a volume %s which served as a snapmirror destination for a previous release of ONTAP.

#### Parameters

**aggrname** (STRING): Name of the aggregate.

**volname** (STRING): Name of the volume.

# openssh events

## openssh.unsupported events

### openssh.unsupported.cipher

**Severity**

ERROR

**Description**

This message occurs when an unsupported cipher is used while attempting an SSH connection.

**Corrective Action**

Attempt to connect with SSH by using a supported cipher.

**Syslog Message**

Attempted an SSH connection using unsupported cipher %s.

**Parameters**

**ciphername** (STRING): Cipher used for the SSH connection.

### openssh.unsupported.hostalgo

**Severity**

ERROR

**Description**

This message occurs when an unsupported host key algorithm is used while attempting an SSH connection.

**Corrective Action**

Attempt to connect with SSH by using a supported host key algorithm.

**Syslog Message**

SSH connection from %s was attempted using an unsupported host key algorithm %s.

**Parameters**

**ipaddress** (STRING): IP address of the remote host that was trying to establish an SSH connection.  
**hostkeyalgorithm** (STRING): Host key algorithm used for the SSH connection attempt.

### openssh.unsupported.kexalgo

**Severity**

ERROR

**Description**

This message occurs when an unsupported key exchange algorithm is used while attempting an SSH connection.

**Corrective Action**

Attempt to connect with SSH by using a supported key exchange algorithm.

**Syslog Message**

Attempted an SSH connection using unsupported key exchange algorithm %s.

**Parameters**

**kexalgorithm** (STRING): Key exchange algorithm used for the SSH connection.

**openssh.unsupported.MAC****Severity**

ERROR

**Description**

This message occurs when an unsupported Message Authentication Code (MAC) is used while attempting an SSH connection.

**Corrective Action**

Attempt to connect with SSH by using a supported MAC.

**Syslog Message**

Attempted an SSH connection using unsupported MAC %s.

**Parameters**

**macname** (STRING): MAC used for the SSH connection.

# osc events

## osc.bucket events

### osc.bucket.notFound

#### Severity

ERROR

#### Description

This message occurs when the object-store bucket or container is not found while trying to access an object.

#### Corrective Action

Verify that the bucket or container has been created on the object-store server and the same name has been correctly configured in ONTAP before accessing the bucket.

#### Syslog Message

Bucket or container "%s" not found on object-store server %s when performing %s operation.

#### Parameters

**bucket** (STRING): Name of the bucket or container.  
**serverHostname** (STRING): Hostname of the object-store server.  
**operation** (STRING): Request name.

## osc.signaturemismatch events

### osc.signatureMismatch

#### Severity

ERROR

#### Description

This message occurs when the request signature sent to the object-store server does not match the signature calculated.

#### Corrective Action

Verify that the secret access key is configured correctly. If it is correct, contact NetApp technical support for assistance.

#### Syslog Message

Object-store %s operation server-calculated request signature does not match the signature sent to object-store server %s for bucket or container "%s" on node %s. Check the keys and signing method.

#### Parameters

**operation** (STRING): Request name.  
**serverHostname** (STRING): Hostname of the object-store server.  
**bucket** (STRING): Name of the object-store bucket.  
**nodeUuid** (STRING): UUID of the node.

# other events

## other.veto events

### other.veto.kmgr.keysmissing

#### Severity

ERROR

#### Description

This message occurs when a giveback or aggregate relocation operation fails but ONTAP® software cannot determine the specific operation type. The specific operation failed due to the unavailability of volume encryption keys for the encrypted volumes of the aggregate on the partner or destination node.

#### Corrective Action

Wait a few minutes, and then try the failed operation again. If the problem persists, run the "security key-manager external restore" command for external key manager, or the "security key-manager onboard sync" command for Onboard Key Manager. Then try failed operation again. The commands retrieve the encryption keys for the encrypted volumes of the aggregate from the key manager.

#### Syslog Message

Relocation or giveback of aggregate %s has failed due to the unavailability of volume encryption keys for the encrypted volumes of the aggregate on the destination (if a relocation operation) or partner node %s (if a giveback operation). Operation code: %d.

#### Parameters

**aggr** (STRING): Name of the aggregate that was not relocated or given back.

**dest\_node** (STRING): Name of the destination node that is missing encryption keys.

**op\_code** (INT): Code of the failed operation

### other.veto.query.timeout

#### Severity

ERROR

#### Description

This message occurs when a query operation on the VEK table times out. The specific operation failed because either a network connection issue occurred or the query could not be completed within the time limit.

#### Corrective Action

Verify that there are no network connectivity issue, and then, if the query still times out, consider overriding the veto.

#### Syslog Message

Relocation or giveback of aggregate "%s" has failed because the restoration status of all of the encryption keys on the destination (if a relocation operation) or partner node (if a giveback operation) "%s" could not be verified within the allotted time.



## Parameters

**aggr** (STRING): Name of the aggregate that was being queried when the timeout occurred.

**dest\_node** (STRING): Name of the destination node.

# pac events

## pac.access events

### pac.access

#### Severity

INFORMATIONAL

#### Description

This message occurs when the access is denied or unsuccessful for the specified source host, destination host, and network interface combination.

#### Corrective Action

If you want to provide access, then modify 'options.(protocol).access' to allow access.

#### Syslog Message

%s: %s from if = %s, source IP = %s, destination IP = %s. Check options %s.access, if desired.

#### Parameters

**protocol1** (STRING): Protocol.  
**status** (STRING): Access status.  
**ifIpAdr** (STRING): Interface address.  
**srcIpAdr** (STRING): Source IP address.  
**dstIpAdr** (STRING): Destination IP address.  
**protocol2** (STRING): Protocol.

## pac.if events

### pac.if.getName

#### Severity

INFORMATIONAL

#### Description

This message occurs when the protocol access control fails to get the interface name with the address.

#### Corrective Action

(None).

#### Syslog Message

%s: getifnamewithaddr() failed with address %s in protocol access control.

#### Parameters

**protocol** (STRING): Protocol.  
**ifAddr** (STRING): Interface IP address in internet standard dotted notation.

# pac.peer events

## pac.peer.getName

### Severity

INFORMATIONAL

### Description

This message occurs when the protocol access control fails to get the peer name.

### Corrective Action

Check Domain Name System (DNS) for possible errors.

### Syslog Message

%s: getpeername() failed in protocol access control with %s.

### Parameters

**protocol** (STRING): Protocol.

**errorMessage** (STRING): Internal error message indicating why the peer name could not be obtained.

# pac.sock events

## pac.sock.getName

### Severity

INFORMATIONAL

### Description

This message occurs when the protocol access control fails to get the socket name.

### Corrective Action

(None).

### Syslog Message

%s: getsockname() failed in protocol access control with %s.

### Parameters

**protocol** (STRING): Protocol.

**errorMessage** (STRING): Internal error message indicating why the socket name could not be obtained.

# passwd events

## passwd.changed events

### passwd.changed

**Severity**

NOTICE

**Description**

The password for a user has been changed.

**Corrective Action**

(None).

**Syslog Message**

passwd for user '%s' changed.

**Parameters**

**username** (STRING): The name of the user whose password has just been changed.

# pci events

## pci.init events

### PCI.init.process.timeWarn

#### Severity

NOTICE

#### Description

This message occurs to warn about the time taken by PCI initialization when PCI configuration is waiting for outstanding background threads that might still be running.

#### Corrective Action

(None).

#### Syslog Message

The PCI initialization process has taken %d seconds.

#### Parameters

**processTime** (INT): Time taken by PCI initialization process.

# pcie events

## pcie.errors events

### pcie.errors

**Severity**

NOTICE

**Description**

This message occurs when correctable PCI Express errors are detected.

**Corrective Action**

(None).

**Syslog Message**

PCI Express correctable error - %s

**Parameters**

**pcie\_errors** (STRING): String contains details about the PCI Express correctable error(s) and which device(s) was/were the source of error(s).

## pcie.link events

### pcie.link.errors

**Severity**

NOTICE

**Description**

This message occurs when PCI Express link errors are detected.

**Corrective Action**

(None).

**Syslog Message**

PCI Express link error - %s

**Parameters**

**pcie\_link\_errors** (STRING): String contains details about the PCI Express link error(s) and which device(s) was/were the source of error(s).

# pcnfsd events

## pcnfsd.auth events

### pcnfsd.auth.failure

**Severity**

ERROR

**Description**

This message occurs when the authorization for a user fails due to a corrupted password entry in `/etc/passwd`.

**Corrective Action**

Restore the password entry for the user.

**Syslog Message**

Corrupted password in `/etc/passwd` file.

**Parameters**

(None).

# pcp events

## pcp.linkstatechangesucceed events

### pcp.linkStateChangeSucceed

#### Severity

NOTICE

#### Description

This message occurs when a change of network port state can now be confirmed after the initial attempt failed due to some race condition. The affected port is now operational.

#### Corrective Action

(None).

#### Syslog Message

Link state change for port %s can now be confirmed after the change %u seconds ago.

#### Parameters

**portname** (STRING): Port name.

**retrySeconds** (INT): Number of seconds that the port state change could not be confirmed.

## pcp.linkstateinprogress events

### pcp.linkStateInProgress

#### Severity

ERROR

#### Description

This message occurs when a change of network port state cannot be confirmed because another configuration task is in progress. The change of the network port state happens when the node is booting up or the administrator changes the network port state manually, via the "network port modify -up-admin [true|false]" command. The affected port might still be operating in the state before the attempted change. While this condition exists, any further configuration or status operations will not be possible.

#### Corrective Action

The link state change request will be retried in the next 5 minutes. If it succeeds, a linkStateChangeSucceed event will be generated. If the issue persists for more than 5 minutes, and the linkStateChangeSucceed event is not generated, there is a configuration task that has not proceeded correctly. The system state will be automatically collected and an AutoSupport® callhome event will be generated. The node from which this event was triggered needs to be rebooted to clear the issue.

#### Syslog Message

Link state change cannot be confirmed on port %s.

#### Parameters

**portname** (STRING): Port name.



# perclientstorepoolthreshold events

## perclientstorepoolthreshold events

### PerClientStorePoolThreshold

#### Severity

NOTICE

#### Description

This message occurs when one of the NFSv4 clients is consuming more resources of a particular store pool type.

#### Corrective Action

(None).

#### Syslog Message

NFSv4 client %s has consumed more resources for store pool type %s.

#### Parameters

**clientIpAddress** (STRING): IP address of the client machine.

**poolName** (STRING): NFSv4 store pool type.

# perf events

## perf.ccma events

### perf.ccma.off

#### Severity

ALERT

#### Description

This message occurs when the performance archiver is not enabled, which makes it difficult for technical support to troubleshoot performance issues on the system.

#### Corrective Action

Enable the performance archiver to track system activity and allow technical support to troubleshoot performance problems by running the "statistics archive config modify -is-enabled true" command. Note that the performance archiver is disabled during takeover and giveback, and in a failure scenario when the archiver loses communication with the Counter Manager Quota Management system; these two scenarios do not resolve by administratively enabling the archiver.

#### Syslog Message

Performance archiver is not enabled for datastore '%s'.

#### Parameters

**datastore** (STRING): Datastore name.

### perf.ccma.startupFailed

#### Severity

ALERT

#### Description

This message occurs when the performance archiver fails to startup.

#### Corrective Action

(None).

#### Syslog Message

Performance archiver failed to startup.

#### Parameters

**reason** (STRING): Reason for failure to startup

### perf.ccma.stop

#### Severity

NOTICE

**Description**

This message occurs when the performance archiver is disabled.

**Corrective Action**

(None).

**Syslog Message**

Performance archiver stopped.

**Parameters**

(None).

# perfstatd events

## perfstatd.dir events

### perfstatd.dir.permission

#### Severity

ERROR

#### Description

This message occurs when the perfstat daemon (perfstatd) startup fails due to improper permissions for the '/mroot/etc/log' directory. This failure impacts the continuous archiving of performance data that technical support uses to troubleshoot performance issues.

#### Corrective Action

Verify that the permissions of the '/mroot/etc/log' directory are 'rwxr-xr-x' by executing the 'fsecurity show /vol/vol0/etc/log' command from the nodelshell. After correcting the permissions for this directory, restart perfstatd by executing the 'spmctl -e -h perfstatd' command from the freebsd shell.

#### Syslog Message

The perfstat daemon startup failed due to improper directory permissions.

#### Parameters

(None).

## perfstatd.disk events

### perfstatd.disk.full

#### Severity

ERROR

#### Description

This message occurs when the perfstat daemon (perfstatd) startup fails due to a full disk. This failure impacts the continuous archiving of performance data that technical support uses to troubleshoot performance issues.

#### Corrective Action

Verify that the '/mroot' file system is full by executing the 'df /mroot' command from the freebsd shell. After clearing some space, restart perfstatd by executing the 'spmctl -e -h perfstatd' command from the freebsd shell.

#### Syslog Message

The perfstat daemon startup failed due to a full disk.

#### Parameters

(None).

# perfstatd.start events

## perfstatd.start.unsuccessful

### Severity

ERROR

### Description

This message occurs when the perfstat daemon (perfstatd) startup fails. This failure impacts the continuous archiving of performance data that technical support uses to troubleshoot performance issues.

### Corrective Action

The startup failure might be due to improper permissions for the '/mroot/etc/log' directory or a full disk. Verify that the permissions of the '/mroot/etc/log' directory are 'rwxr-xr-x' by executing the 'fsecurity show /vol/vol0/etc/log' command from the nodeshell. Verify that the '/mroot' file system is not full. After fixing the cause of the failure, restart perfstatd by executing the 'spmctl -e -h perfstatd' command from the freebsd shell.

### Syslog Message

The perfstat daemon startup failed.

### Parameters

(None).

# pha events

## pha.add events

### pha.add.dev.limitExceeded

#### Severity

ERROR

#### Description

This message occurs when an attempt to add a passthrough host adapter (PHA) disk to a virtual disk adapter fails because the supplied disk number exceeds the maximum supported disk ID.

#### Corrective Action

Reduce the number of devices connected to Data ONTAP® and retry this operation.

#### Syslog Message

Failed to add a device with unit number %d to virtual disk adapter. Reason: Supplied disk number exceeds the max number of supported disks.

#### Parameters

**disk\_num** (INT): Unit number for the rejected disk.

## pha.dev events

### pha.dev.get.capacity.failed

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software fails to get a drive's capacity after an extended wait. As a result, cannot add the drive.

#### Corrective Action

Ensure that the hypervisor's device is still connected and accessible by ONTAP.

#### Syslog Message

Failed to get drive capacity for device %s%d after %d seconds.

#### Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

**wait\_time** (INT): Number of seconds waited for capacity information.

### pha.dev.get.uuid.bioReadErr

#### Severity

ERROR

## Description

This message occurs when block read operation for the passthrough host adapter (PHA) disk header returns an error. As a result, Data ONTAP® is unable to add or manage this device.

## Corrective Action

Ensure that the device is still connected to and accessible by Data ONTAP® and retry the operation.

## Syslog Message

Failed to get UUID for device %s%d. Reason: Block read operation returned an error: 0x%x.

## Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

**bio\_error** (INTHEX): Error code returned by block read operation.

## pha.dev.get.uuid.devNotAdded

### Severity

ERROR

### Description

This message occurs when Data ONTAP® fails to get disk header after multiple retries and finally give up. As a result, Data ONTAP® is unable to add the disk.

### Corrective Action

Ensure that the device is still connected to and accessible by Data ONTAP® and retry the operation.

### Syslog Message

Failed to get disk header for device %s%d. Number of bioerr %d Retries %d.

### Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

**bioerr\_cnt** (INT): Number of BIO errors.

**retry\_cnt** (INT): Number of retries.

## pha.dev.get.uuid.devNotExist

### Severity

ERROR

### Description

This message occurs when Data ONTAP® fails to retrieve the the descriptor for the passthrough host adapter (PHA) device. As a result, Data ONTAP® fails to read the header on this device and is unable to add or manage it.

### Corrective Action

Ensure that the device is still connected to and accessible by Data ONTAP® and retry the operation.

### Syslog Message

Failed to retrieve the descriptor for device %s%d. Reason: %s.

## Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

**error** (STRING): Error message.

## pha.dev.header.err

### Severity

ERROR

### Description

This message occurs when a disk header operation encounters a device error. As a result, ONTAP® might be unable to add or manage this device.

### Corrective Action

Ensure that the device is still connected to and accessible by ONTAP. If the device is reported as healthy by the hypervisor, contact NetApp technical support to open a support case to investigate the underlying cause.

### Syslog Message

Functional operation "%s" failed for device %s%d. Error number: 0x%x, CAM status: 0x%x, SCSI (status: 0x%x, sense key: x%x, sense code: 0x%x, additional sense code: 0x%x, additional sense code qualifier: 0x%x).

## Parameters

**operation** (STRING): Name of the functional operation.

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

**bio\_error** (INTEHEX): Error code returned by the device.

**cam\_status** (INTEHEX): Common Access Method (CAM) status code returned by the device.

**scsi\_status** (INTEHEX): SCSI status code returned by the device, if available.

**error\_code** (INTEHEX): SCSI error code returned by the device, if available.

**sense\_key** (INTEHEX): SCSI sense key returned by the device, if available.

**asc** (INTEHEX): SCSI additional sense code returned by the device, if available.

**ascq** (INTEHEX): SCSI additional sense code qualifier returned by the device, if available.

## pha.dev.uuidWriteHeaderErr

### Severity

ERROR

### Description

This message occurs when Data ONTAP® fails to write the header to the passthrough host adapter (PHA) device. As a result, Data ONTAP® is unable to add or manage this device.

### Corrective Action

Ensure that the device is still connected to and accessible by Data ONTAP® and retry the operation.

### Syslog Message

Failed to write header to the device %s%d.



## Parameters

**devname** (STRING): Name of the device.

**unit** (INT): Unit number of the device.

# pha.disk events

## pha.disk.type.changed

### Severity

ERROR

### Description

This message occurs when the backing disk type for an existing disk was changed from the cloud API.

### Corrective Action

Check whether the disk type of the failed disk was changed. Change the backing disk type to the original value.

### Syslog Message

Backing disk type changed for Disk: %s, Volume: %s. Error: %s.

### Parameters

**d\_name** (STRING): Disk name

**d\_id** (STRING): Disk ID

**errmsg** (STRING): Indicating the stage at which the disk lookup failed.

## pha.disk.type.lookup.failed

### Severity

ERROR

### Description

This message occurs when the backing disk type lookup fails.

### Corrective Action

Check that the disk type of the failed disk is supported. Remove the affected disk.

### Syslog Message

Failed to recognize backing disk type for Disk: %s, Volume: %s. Error: %s.

### Parameters

**d\_name** (STRING): Disk name.

**d\_id** (STRING): Disk ID

**errmsg** (STRING): Indicating the stage at which the disk lookup failed.

# pha.obj events

## pha.obj.blobGone

### Severity

NOTICE

### Description

This message occurs when ONTAP® detects that a blob has been deleted from a container.

### Corrective Action

If the identified blob was not intentionally deleted, attempt to recover it with the object store provider.

### Syslog Message

Blob "%s" is no longer seen in object store, "%s".

### Parameters

**blob** (STRING): Object name.

**configName** (STRING): Object store configuration name.

## pha.obj.scanFailure

### Severity

ERROR

### Description

This message occurs when ONTAP® cannot get a list of objects from the object store server.

### Corrective Action

If the error text indicates an authentication problem, verify that the stored credentials are correct by using the "system node virtual-machine disk-object-store modify" command. If the error text describes a connection problem, determine the IP address for the object store server. If the address reported in this event does not match, examine the DNS cache for errors by using the "system node virtual-machine disk-object-store dns-cache show" command. If the address does match, check the status of both the object store server and the network.

### Syslog Message

ONTAP could not get a list of objects from the object store, "%s", at "%s", using IP address "%s" with error: %s.

### Parameters

**configName** (STRING): Object store configuration name.

**server** (STRING): Object store server name.

**ipAddress** (STRING): Object store IP address.

**error** (STRING): Error string.

## pha.obj.scanSuccess

### Severity

NOTICE

### Description

This message occurs when ONTAP® succeeds at retrieving a list of objects from the object store server after reporting a failure.

**Corrective Action**

(None).

**Syslog Message**

ONTAP successfully retrieved a list of objects from the object store, "%s".

**Parameters**

**configName** (STRING): Object store configuration name.

**pha.obj.storeAvailable****Severity**

NOTICE

**Description**

This message occurs when ONTAP® can communicate with the object store server.

**Corrective Action**

(None).

**Syslog Message**

ONTAP successfully communicated with the object store, "%s".

**Parameters**

**configName** (STRING): Object store configuration name.

**pha.obj.storeCreateFailed****Severity**

ERROR

**Description**

This message occurs when ONTAP® software cannot create the object store configuration.

**Corrective Action**

Verify that the values provided for the create operation are correct. ONTAP supports either the HTTPS or HTTP protocol and not a mixed configuration. The value for the port should be the same as the first object store's (object\_store0) port. Verify the first object store's port by using the "system node virtual-machine disk-object-store show -object-store-name object\_store0 -node node\_name" command, which will display the port as "Port Number of the Object Store".

**Syslog Message**

ONTAP could not create the object store, "%s", due to reason: %s.

**Parameters**

**configName** (STRING): Object store configuration name.

**reason** (STRING): An object store create operation was attempted with a port value that wasn't equal to that of object\_store0.

## pha.obj.storeDeleted

### Severity

NOTICE

### Description

This message occurs when ONTAP® finishes deleting the connection to the object store server.

### Corrective Action

(None).

### Syslog Message

ONTAP has removed the object store configuration, "%s".

### Parameters

**configName** (STRING): Object store configuration name.

## pha.obj.storeDeleting

### Severity

NOTICE

### Description

This message occurs when ONTAP® begins to tear down the connection to the object store server.

### Corrective Action

(None).

### Syslog Message

ONTAP began removing the object store configuration, "%s".

### Parameters

**configName** (STRING): Object store configuration name.

## pha.obj.storeUnavailable

### Severity

EMERGENCY

### Description

This message occurs when ONTAP® cannot communicate with the object store server.

### Corrective Action

Verify that the configuration of the object store, including the login credentials, is still accurate by using the "system node virtual-machine disk-object-store show" command. Check the network status for connectivity errors.

### Syslog Message

ONTAP could not communicate with the object store, "%s".

## Parameters

**configName** (STRING): Object store configuration name.

## pha.obj.storeUpdateFailed

### Severity

ERROR

### Description

This message occurs when ONTAP® cannot update the object store configuration.

### Corrective Action

Verify that the updates are correct. Then create a new object store configuration with the updated fields. Delete the old configuration after ONTAP detects the new disk paths.

### Syslog Message

ONTAP could not update the object store configuration, "%s", due to reason: %s.

## Parameters

**configName** (STRING): Object store configuration name.

**reason** (STRING): Reason for the failure.

## pha.obj.throttle

### Severity

NOTICE

### Description

This message occurs when the passthrough host adapter object (PHA Obj) throttled on the disks due to receiving the OSC\_CONNECTION\_UNAVAILABLE, OSC\_TIMEOUT or AZURE\_SERVER\_BUSY error from OSC.

### Corrective Action

(None).

### Syslog Message

PHA Obj throttled %llu times on the blob %s in the container %s due to the "%s".

## Parameters

**counter** (LONGINT): The throttle counter.

**blobName** (STRING): The throttling blob name.

**containerName** (STRING): The throttling container name.

**throttleReason** (STRING): The error which triggered throttling.

## pha.obj.unknownTier

### Severity

ERROR

## Description

This message occurs when ONTAP® does not recognize the tier of a page blob.

## Corrective Action

Change the blob tier to one that ONTAP will recognize: P10, P15, P20, P30, P40, P50, or P60.

## Syslog Message

ONTAP did not recognize tier "%s" of blob "%s", on object store "%s". Assuming Standard Storage Tier.

## Parameters

**tier** (STRING): Blob tier.

**blob** (STRING): Blob name.

**configName** (STRING): Object store configuration name.

# pha.vsa events

## pha.vsa.uuid.too.long

### Severity

ALERT

### Description

This message occurs when ONTAP® software detects a device with a device descriptor length greater than 40 characters. Such devices are not presented to ONTAP.

### Corrective Action

Detach the device from all ONTAP virtual machines, and then delete it. Create a similar device with a descriptor length less than or equal to 40 characters, and then attach it to the ONTAP virtual machines.

### Syslog Message

The device descriptor %s is too long and was not presented to ONTAP.

### Parameters

**Descriptor** (STRING): Descriptor of the device.

# phf events

## phf.manage events

### phf.manage.add.upcall

#### Severity

NOTICE

#### Description

This message occurs when peripheral host bus adapter freebsd (phf) detects a device to be added to the system.

#### Corrective Action

(None).

#### Syslog Message

Added unit %d type %d.

#### Parameters

**unit** (INT): Unit number.

**pman\_disk\_type** (INT): Disk type.

# platform events

## platform.chsismobonamediff events

### platform.ChsisMoboNameDiff

#### Severity

ERROR

#### Description

This message occurs when the product name on the chassis is different from the one on the controller.

#### Corrective Action

Place the controller in a compatible chassis.

#### Syslog Message

The product names on the chassis (%s) and the controller (%s) are different. Contact NetApp technical support to resolve this.

#### Parameters

**prodNameChsis** (STRING): Product name on the chassis.  
**prodNameCtrlr** (STRING): Product name on the controller.

## platform.dimmconfigerror events

### platform.dimmConfigError

#### Severity

ALERT

#### Description

This message occurs when the DIMM configuration of the controller is different from the configuration supported by this model. This often occurs when an incorrect DIMM memory size is used in a DIMM slot.

#### Corrective Action

Place correctly sized and supported DIMMs in the designated slots. For replacement instructions, search the NetApp technical support web site using the keywords "replacing", "DIMM", and "system".

#### Syslog Message

Memory DIMM configuration mismatch in slots %s.

#### Parameters

**slots** (STRING): List of slots with the incorrect DIMM.

## platform.insufficientmemory events



## platform.insufficientMemory

### Severity

EMERGENCY

### Description

This message occurs when there is not enough system memory for the system to function, likely due to slots that contain incorrect, failed, or missing DIMMs.

### Corrective Action

Replace the incorrect, failed, or missing DIMMs.

### Syslog Message

Not enough system memory for the system to function. Detected %d GB, expected %d GB. Check DIMMs slots %s.

### Parameters

**mem\_detected** (INT): Amount of system memory that was detected.

**mem\_expected** (INT): Amount of system memory that this model supports.

**slots** (STRING): List of slots with incorrect, failed, or missing DIMMs.

## platform.nopartnerpresent events

### platform.noPartnerPresent

### Severity

NOTICE

### Description

This message occurs when the HA partner's presence cannot be read via the GPIO ports on the controller.

### Corrective Action

(None).

### Syslog Message

Controller partner's presence could not be read via GPIO. Read partner ID: %d.

### Parameters

**partner\_id** (INT): The partner ID that was read from the GPIO pins.

## platform.reducedmemory events

### platform.reducedMemory

### Severity

ALERT

### Description

This message occurs when system memory is less than expected for this model, likely due to slots that contain incorrect, failed, or missing DIMMs.

**Corrective Action**

Replace the incorrect, failed, or missing DIMMs.

**Syslog Message**

System memory (%d GB) is less than expected (%d GB). Check DIMMs slots %s.

**Parameters**

**mem\_detected** (INT): Amount of system memory that was detected.

**mem\_expected** (INT): Amount of system memory that this model supports.

**slots** (STRING): List of slots with incorrect, failed, or missing DIMMs.

# pn events

## pn.driver events

### pn.driver.send.latency.high

**Severity**

NOTICE

**Description**

This message occurs when the pass-through network layer detects driver send latencies that are too high, indicating a noisy neighbor issue in cloud providers that is causing an I/O latency spike.

**Corrective Action**

(None).

**Syslog Message**

Pass-through network layer is reporting an average driver send latency of %d microseconds, which might cause an I/O latency spike.

**Parameters**

**latency** (INT): Average driver send latency in microseconds.

## pn.mvia events

### pn.mvia.freeq.empty

**Severity**

NOTICE

**Description**

This message occurs when the pass-through network layer depletes its supply of buffers for sending traffic, which could cause an Interconnect timeout.

**Corrective Action**

(None).

**Syslog Message**

Pass-through network layer free queue empty for %d msec, which might cause Interconnect timeout.

**Parameters**

**time** (INT): Duration of depletion in msec.

### pn.mvia.nic.reset

**Severity**

NOTICE

**Description**

This message occurs when the passthrough network layer detects network interface card (NIC) state change caused by an NIC reset.

**Corrective Action**

(None).

**Syslog Message**

Passthrough network layer detected NIC %s has changed its operational state. New state: %s.

**Parameters**

**nic\_name** (STRING): NIC name.

**nic\_state\_change** (STRING): NIC state change.

**pn.mvia.send.starvation****Severity**

NOTICE

**Description**

This message occurs when the pass-through network layer takes too long to execute the logic for sending traffic, which could cause an Interconnect timeout.

**Corrective Action**

(None).

**Syslog Message**

Pass-through network layer "send" thread not run for %d msec, which might cause Interconnect timeout.

**Parameters**

**time** (INT): How long the pass-through network layer has not executed the logic for sending traffic.

# probedone events

## probedone.inquiry events

### probedone.inquiry.failed

#### Severity

ERROR

#### Description

This message occurs when the inquiry data fails the comparison at the domain validation (DV) step.

#### Corrective Action

Check whether the device is connected correctly.

#### Syslog Message

In path %s, inquiry data fails comparison at the domain validation DV%d step of protocol type %d.

#### Parameters

**path** (STRING): Device path.

**inquirydata** (INT): Comparison between the inquiry data failure and the domain validation, and whether the probe inquiry domain validation is type 1 or 2.

**protocol\_type** (INT): Protocol type.

# protocol events

## protocol.exportcheckfailed events

### protocol.exportCheckFailed

#### Severity

ERROR

#### Description

This message occurs when there is an export check failure due to, either export-policy configuration not allowing access, or a non-existing filesystem/junction access attempt.

#### Corrective Action

Check the export policy configuration or filesystem name depending on the reason due to which export access check failed.

#### Syslog Message

Export access check failed on Vserver %s for client %s, trying to access volume %s, having policy %s, via protocol %s. Reason: %s

#### Parameters

**vserver** (STRING): Name of the Vserver where the issue occurs.

**clientIP** (STRING): Client for which the export access check failure happened.

**volumeName** (STRING): Volume name for which export access check failed.

**policyName** (STRING): Name of the export policy from which export access check failed.

**protocol** (STRING): Access protocol used by client for filesystem access.

**reason** (STRING): Reason for which export access check failed.

# prov events

## prov.rollback events

### prov.rollback.fail.cmd

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software cannot complete a rollback of a provisioning operation sequence that failed. Depending on the action that failed, objects might need to be created, deleted, or modified to bring the system back to a consistent state.

#### Corrective Action

Resolve the error and attempt the failed command again. If the problem persists, contact NetApp technical support to resolve the rollback failure. It is possible that other related rollback actions have also failed. Check for other 'prov.rollback.fail.\*' messages with similar time stamps, and perform the suggested corrective actions in the same order that the messages occurred.

#### Syslog Message

Failed to roll back an action performed by the provisioning engine. Command: "%s". Reason: %s.

#### Parameters

**failed\_command** (STRING): Command that failed.

**error** (STRING): Error that the provisioning engine encountered.

### prov.rollback.fail.obj

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software cannot complete a rollback of a provisioning operation sequence that failed. Depending on the action that failed, objects might need to be created, deleted, or modified to bring the system back to a consistent state.

#### Corrective Action

Contact NetApp technical support to resolve the rollback failure. It is possible that other related rollback actions have also failed. Check for other 'prov.rollback.fail.\*' messages with similar time stamps, and perform the suggested corrective actions in the same order that the messages occurred.

#### Syslog Message

Failed to roll back an action performed by the provisioning engine. Action: "%s". Object: "%s". Fields: "%s". Reason: %s.

#### Parameters

**action** (STRING): Attempted action.

**object** (STRING): Type of object.

**fields** (STRING): Fields of the object.

**error** (STRING): Error that the provisioning engine encountered.

## **prov.rollback.fail.rest**

### **Severity**

ERROR

### **Description**

This message occurs when ONTAP® software cannot complete a rollback of a provisioning operation sequence that failed. Depending on the action that failed, objects might need to be created, deleted, or modified to bring the system back to a consistent state.

### **Corrective Action**

Resolve the error and attempt the failed REST API again. If the problem persists, contact NetApp technical support to resolve the rollback failure. It is possible that other related rollback actions have also failed. Check for other 'prov.rollback.fail.\*' messages with similar time stamps, and perform the suggested corrective actions in the same order that the messages occurred.

### **Syslog Message**

Failed to roll back an action performed by the provisioning engine. API: curl -X %s 'https://<mgmt-ip>%s' -H 'accept: application/hal+json' -d '%s' Reason: %s.

### **Parameters**

**failed\_uri** (STRING): REST API resource that failed.

**failed\_body** (STRING): REST API JSON body that failed.

**failed\_method** (STRING): REST API method that failed.

**error** (STRING): Error that the provisioning engine encountered.



# pubsub events

## pubsub.kernel events

### pubsub.kernel.flow.control.enabled

#### Deprecated

This event is deprecated as of 9.10.0.

#### Severity

ERROR

#### Description

This message will not be published as flow control is enabled on kernel pubsub.

#### Corrective Action

Check the network connectivity to the broker and ensure the broker is running properly.

#### Syslog Message

Kernel pubsub failed to send message with topic %s, object %s, object UUID %s and operation %s. Flow control is enabled.

#### Parameters

**topic** (STRING): Topic of the message.  
**object** (STRING): Object of the message.  
**object\_uuid** (STRING): Object UUID of the message.  
**operation** (STRING): Operation of the message.

### pubsub.kernel.send.msg.failed

#### Deprecated

This event is deprecated as of 9.10.0.

#### Severity

ERROR

#### Description

This message occurs when kernel pubsub fails to send a message to the broker.

#### Corrective Action

Check the network connectivity to the broker and ensure the broker is running properly.

#### Syslog Message

Kernel pubsub failed to send message with topic %s, object %s, object UUID %s, operation %s, and error code %d.

#### Parameters

**topic** (STRING): Topic of the message.  
**object** (STRING): Object of the message.  
**object\_uuid** (STRING): Object UUID of the message.

**operation** (STRING): Operation of the message.

**error\_code** (INT): Error code associated with the failure.

# pvif events

## pvif.badmtu events

### pvif.badMtu

#### Severity

ERROR

#### Description

This message occurs when the interface group (ifgrp) driver fails to update the MTU size on a member port because it does not support the requested configuration.

#### Corrective Action

Use the "ifgrp remove-port" command to remove the member port that does not support the requested MTU size, or use the "broadcast-domain remove-ports" and "broadcast-domain add-ports" commands to assign the ifgrp to a broadcast domain whose MTU are supported on all member ports.

#### Syslog Message

ifgrp %s: unable to update MTU size to %d on %s.

#### Parameters

**ifgrpName** (STRING): Name of the ifgrp.

**mtuSize** (INT): Requested MTU size.

**portName** (STRING): Name of the member port.

## pvif.lacp events

### pvif.lacp.outofsync

#### Severity

ALERT

#### Description

This message occurs when a link is operationally up and cannot be aggregated into its configured ifgrp within the time limit. This might occur due to misconfiguration, physical link error, or implementation limits on the number of links allowed.

#### Corrective Action

Verify that the switch configuration is correct, and that the physical links are operating without error.

#### Syslog Message

ifgrp %s, port %s cannot synchronize with the LACP aggregate.

#### Parameters

**vifName** (STRING): Name of the ifgrp.

**linkName** (STRING): Name of the link.

# qos events

## qos.bgtask events

### qos.BgTask.Throttle

#### Severity

NOTICE

#### Description

This message occurs when the Quality of Service (QoS) subsystem starts throttling requests to help critical background tasks meet their min-throughput and their service-level objectives (SLOs).

#### Corrective Action

(None).

#### Syslog Message

QoS throttle limit for background task: "%s" with id: "%u" is %u IOPS.

#### Parameters

**bg\_task\_name** (STRING): Name of the Background Task that requested QoS to throttle requests.

**bg\_task\_id** (INT): ID of the Background Task that requested QoS to throttle requests.

**limit** (INT): The limit value enforced in terms of IOPS and a block size of 4K.

## qos.mgmt events

### qos.mgmt.operation.failed

#### Severity

ERROR

#### Description

This message indicates that cluster quality of service (QoS) cannot propagate settings to the classification and control components on the specified node. This might be a transient condition, or it might indicate insufficient resources to complete the operation. QoS workload classification and control might behave incorrectly on this node.

#### Corrective Action

Inspect the management gateway process log for errors indicating why the command failed. If this behavior persists, contact NetApp technical support.

#### Syslog Message

QoS configuration update failed on node %s: %s

#### Parameters

**node** (STRING): Node on which operation failed.

**error** (STRING): Error description.

# qos.min events

## qos.min.tput.not.honoured.for.sync

### Severity

NOTICE

### Description

This message occurs when minimum throughput value is not honoured if adaptive QoS is set on a SnapMirror Synchronous protected storage object.

### Corrective Action

(None).

### Syslog Message

Adaptive QoS policy "%s" set on SnapMirror Synchronous volume "%s" or any of it's storage object has a min-throughput value set. SnapMirror Synchronous does not support min-throughput hence min-throughput setting will not be honoured from the policy, only max-throughput setting will be honoured.

### Parameters

**policy** (STRING): Policy group name.

**volume** (STRING): Volume name.

# qos.mintp events

## qos.mintp.high.intrconct.lat

### Severity

NOTICE

### Description

This message occurs when the Quality of Service (QoS) subsystem notices high latency in the cluster interconnect for workloads associated with minimum-throughput policy groups. As a result, workloads might not be able to meet their minimum-throughput values.

### Corrective Action

Consider using the "qos statistics" CLI command to monitor the workloads and policy groups that have high cluster interconnect latency and also do not meet their minimum-throughput Service Level Objectives (SLOs). Either avoid the interconnect by moving workloads or LIFs, or reduce contention for the interconnect by applying QoS maximum throughput constraints on interfering workloads.

### Syslog Message

QoS has noticed high cluster interconnect latency. Some workloads and policy groups might not meet their specified minimum-throughput.

### Parameters

(None).

# qos.monitor events

## qos.monitor.memory.abated

### Severity

NOTICE

### Description

This message occurs when the Quality of Service (QoS) subsystem's dynamic memory is no longer at its limit for the current platform hardware. All QoS features will operate as normal.

### Corrective Action

(None).

### Syslog Message

QoS dynamic memory is no longer at its limit. All QoS features will operate as normal.

### Parameters

**subject** (STRING): QoS Memory message title.

**count** (INT): Number of times this has happened.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## qos.monitor.memory.maxed

### Severity

ERROR

### Description

This message occurs when the Quality of Service (QoS) subsystem's dynamic memory hits its limit for the current platform hardware. As a result, some QoS features might operate in a limited capacity.

### Corrective Action

Some active workloads or streams should be deleted to free up memory. Use the 'statistics show -object workload -counter ops' command to determine which workloads are active. Active workloads will show non-zero ops. Then use the 'workload delete [workload name]' command multiple times to remove particular workloads. Alternatively, use the 'stream delete -workload [workload name] \*' command to delete the associated streams from the active workload.

### Syslog Message

QoS dynamic memory has reached its limit. Some QoS features might operate in a limited capacity.

### Parameters

**subject** (STRING): QoS Monitor message title.

**count** (INT): Number of times this has happened.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## qos.onnix events

### qos.onnix.classf.failure

#### Severity

ERROR

#### Description

This message occurs when the classification subsystem for Quality of Service (QoS) fails. Any limits that are configured for volumes will not be applied.

#### Corrective Action

Contact Contact NetApp technical support. for assistance.

#### Syslog Message

QoS classification failed. Any limits that are configured for volumes will not be applied.

#### Parameters

(None).

## qos.qtree events

### qos.qtree.fh.error.report

#### Severity

ERROR

#### Description

This message occurs when the Quality of Service (QoS) subsystem recognizes that an older version of a filehandle is still being used to access a qtree associated with a QoS policy group. Use of this filehandle will impact the ability of the qtree workload to meet its service-level objectives (SLOs).

#### Corrective Action

Use the "qos workload show -wid workloadID" command to show the affected qtree. Use the "network connections active show -local-port 2049" command to show all of the active NFS protocol connections to each node. Verify that, after a QoS policy group was associated with the qtree, each client did in fact remount all of its mount points with the qtree. If any client did not remount, first determine which Vserver exported the mount points by using the "volume qtree show -fields qtree,vserver" command. Next stop all applications that are using the qtree and unmount all exports from the Vserver in question. Finally, remount the exports and restart the applications.

#### Syslog Message

QoS qtree violation for qtree with workload ID = %u.

#### Parameters

**wid** (INT): The QoS workload ID of the qtree.

## qos.viodet events

## qos.VioDet.maxThrottle

### Severity

NOTICE

### Description

This message occurs when the Quality of Service (QoS) subsystem has already throttled requests severely to help workloads meet their min-throughput and their service-level objectives (SLOs) and will not throttle any further. Possible reasons for severe throttling could be that there are bully workloads or insufficient resources, which means the system cannot support the specified targets in the SLOs.

### Corrective Action

Collect ONTAP statistics counters and use 'qos statistics' CLI command on the node to monitor the workload and policy group to find the workloads that do not meet their min-throughput SLOs. The cpu utilization and aggregate utilization can be used to help identify if there is sufficient load or if the system is over-provisioned. First validate if the policy group client generates sufficient load. If the system is over provisioned, identify the critical resources (cpu, memory, network, disk, etc.) that could lead to the violation, or identify the hot spot in a clustered system and move workloads around to avoid the hot spot or limit bully workloads by setting a QoS limit to reduce the contention.

### Syslog Message

QoS violation limit for %s %s is %u IOPS.

### Parameters

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

**vdlimit** (INT): Current limit value in terms of IOPS and a block size of 4K.

## qos.VioDet.Mintput.Throttle

### Severity

ALERT

### Description

This message occurs when the Quality of Service (QoS) subsystem starts throttling requests to help workloads meet their min-throughput and their service-level objectives (SLOs). Possible reasons for throttling could be that there are bully workloads or insufficient resources, which means the system cannot support the specified targets in the SLOs.

### Corrective Action

Collect ONTAP statistics counters and use 'qos statistics' CLI command on the node to monitor the workload and policy group to find the workloads that do not meet their min-throughput SLOs. The CPU utilization and aggregate utilization can be used to help identify if there is sufficient load or if the system is over-provisioned. Firstly validate if the policy group client generates sufficient load. If the system is over provisioned, identify the critical resources (CPU, memory, network, disk, etc.) that could lead to the violation, or identify the hot spot in a clustered system and move workloads around to avoid the hot spot. Or limit bully workloads by setting a QoS limit to reduce the contention.

### Syslog Message

QoS min-throughput feature has started throttling requests for node: "%s" with uuid: "%s".



## Parameters

**object\_type** (STRING): Type of resource object under notification. The object\_type will always be NODE.  
**object\_uuid** (STRING): Universal unique identifier (UUID) of the resource object. The UUID is of the node.

## qos.VioDet.noMemory

### Severity

ERROR

### Description

This message occurs when the Quality-of-Service (QoS) process cannot allocate memory required for the Violation Detection subsystem (likely there is not enough system memory available). As a result workloads will not meet their min-throughput and their service-level objectives (SLOs).

### Corrective Action

(None).

### Syslog Message

QoS dynamic memory has reached its limit. QoS min-throughput might operate in a limited capacity.

## Parameters

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## qos.VioDet.publishFail

### Severity

ERROR

### Description

This message occurs when the Quality-of-Service (QoS) process cannot communicate throttle limits due to an internal error. As a result workloads might not meet their min-throughput and their service-level objectives (SLOs).

### Corrective Action

Contact NetApp technical support.

### Syslog Message

QoS subsystem cannot publish throttle limits.

## Parameters

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

## qos.violation events

## qos.violation.report

### Deprecated

Deprecated as of version 9.4.0 since the QoS violation detection functionality has been deprecated.

### Severity

NOTICE

### Description

This message occurs every hour to provide a snapshot of violation detection status for the Quality of Service (QoS) policy groups. The snapshot shows the states of those policy groups that either cannot meet their service-level objectives (SLOs), or are impacted by (deactivation) actions taken by QoS to remedy the situations. Possible violation states include `LOW_OFFERED`, which means the offered load for a policy group is too low, or `INSUFFICIENT_RESOURCES`, which means the system can not support the specified targets in the SLO. The deactivation states include `ACTIVATED`, which means this policy group is not impacted by deactivation, or `DEACTIVATED`, which means this policy group is deactivated to support other violated policy group.

### Corrective Action

After observing a violation, one should collect ONTAP statistics counters and use 'qos statistics' CLI command to monitor the violated policy group. Increase the client load if a low offered violation is present and the policy group client does not generate sufficient load. For other violations, identify the critical resources (cpu, memory, network, disk, etc.) that could lead to the violation, or identify the hot spot in a clustered system. The solution to the former is to add more resources; while the latter is to move workloads around to avoid the hot spot.

### Syslog Message

QoS violation state for policy group %s is (pgid= %u violation= %s deactivation= %s).

### Parameters

**name** (STRING): The QoS policy group name.

**id** (INT): The QoS policy group identifier.

**vdstate** (STRING): The violation state for a QoS policy group.

**dastate** (STRING): The deactivation state for a QoS policy group.

# qtree events

## qtree.delete events

### qtree.delete.failed

#### Severity

ERROR

#### Description

This message occurs when the delete operation on a qtree fails.

#### Corrective Action

Check and resolve the failure reason and attempt the qtree deletion again.

#### Syslog Message

Failed to delete qtree "%s" in volume "%s" on Vserver "%s". Error: "%s" (%d).

#### Parameters

**qtree** (STRING): Name of the qtree.

**volume** (STRING): Name of the volume on which the qtree resides.

**vserver** (STRING): Name of the Vserver on which the volume resides.

**reason** (STRING): Reason for the failure.

**errorCode** (INT): Error number of the failure.

### qtree.delete.start

#### Severity

NOTICE

#### Description

This message occurs when a qtree is being deleted. A subsequent event will indicate the success or failure of the operation.

#### Corrective Action

(None).

#### Syslog Message

Deleting qtree "%s" in volume "%s" on Vserver "%s".

#### Parameters

**qtree** (STRING): Name of the qtree.

**volume** (STRING): Name of the volume on which the qtree resides.

**vserver** (STRING): Name of the Vserver on which the volume resides.

### qtree.delete.success

#### Severity

NOTICE

**Description**

This message occurs when the delete operation on a qtree succeeds.

**Corrective Action**

(None).

**Syslog Message**

Successfully deleted qtree "%s" in volume "%s" on Vserver "%s".

**Parameters**

**qtree** (STRING): Name of the qtree.

**volume** (STRING): Name of the volume on which the qtree resides.

**vserver** (STRING): Name of the Vserver on which the volume resides.

# query events

## query.execution events

### query.execution.console

**Severity**

NOTICE

**Description**

This message occurs when the dotsql command-line program processes an SQL query.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**query** (STRING): Query that was executed.

### query.execution.failed

**Severity**

NOTICE

**Description**

This message occurs when a SQL query of the system fails for any reason.

**Corrective Action**

This type of failure can occur legitimately during the course of normal system operation. In cases of unexpected failure, the requestor provides diagnostic advice.

**Syslog Message**

(None).

**Parameters**

**query** (STRING): Query that was executed.

**error** (STRING): Error message that occurred.

# quota events

## quota.inconsistent events

### quota.inconsistent

#### Severity

ERROR

#### Description

This message occurs when the quota database is corrupted.

#### Corrective Action

Use the 'quota off' command followed by the 'quota on' command to clear the corruption and reenables quotas on the volume.

#### Syslog Message

Quota database on %s%s%s is corrupted. Disabling quotas until the corruption is fixed.

#### Parameters

**vol\_name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

## quota.intr events

### quota.intr.boot.fschk

#### Severity

INFORMATIONAL

#### Description

This message occurs when previous quota initialization was interrupted due to a system reboot or file system consistency checking.

#### Corrective Action

(None).

#### Syslog Message

WAFL: Reboot or file system consistency checking interrupted quota initialization on %s %s%s%s%s.

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

# quota.parse events

## quota.parse.error

### Severity

ERROR

### Description

This message occurs when the system detects an error while scanning the quota rules. The parser would be skipping this quota definition.

### Corrective Action

Fix the error in quota rules definition using the CLI commands "quota policy rule modify" or "quota policy rule delete".

### Syslog Message

%s

### Parameters

**errmsg** (STRING): Message describing the error in detail.

# quota.push events

## quota.push.rules.complete

### Severity

NOTICE

### Description

This message occurs when the system has completed pushing rules into the quota database.

### Corrective Action

(None).

### Syslog Message

Quota rules push completed on %s %s%s%s%s.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

## quota.push.rules.start

### Severity

NOTICE

## Description

This message occurs when the system starts pushing rules into the quota database.

## Corrective Action

(None).

## Syslog Message

Quota rules push started on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

# quota.reinit events

## quota.reinit.required

### Severity

ERROR

### Description

This message occurs when an unplanned MetroCluster(tm) switchover happens. If a quota operation was running on a volume located on a disaster site during the unplanned switchover, the quota database could be in an incomplete state and quotas need to be reinitialized on that volume.

### Corrective Action

After the unplanned switchover is complete, use the "quota off" CLI command followed by the "quota on" CLI command to reinitialize quotas for the indicated volume.

### Syslog Message

Quota reinitialization is required on %s%s%s due to an unplanned switchover.

### Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

# quota.resize events

## quota.resize.start

### Severity

NOTICE



**Description**

This message occurs when quota resize starts.

**Corrective Action**

None.

**Syslog Message**

%s

**Parameters**

**noticestring** (STRING): Quota resize started.

**quota.resize.stop****Severity**

NOTICE

**Description**

This message occurs when quota resize stops or finishes.

**Corrective Action**

None.

**Syslog Message**

%s

**Parameters**

**noticestring** (STRING): Quota resize stopped or completed. It also reports the duration for the quota resize.

**quota.scan events****quota.scan.abort****Severity**

ERROR

**Description**

This message occurs after the administrator uses the "volume quota on" command, when the quota scanner aborts without finishing due to an error condition. Any quotas configured on the volume will not be tracked or enforced.

**Corrective Action**

If the volume quota status is not displayed as "off", turn quotas off by executing the "volume quota off" command, and then retry the "volume quota on" command. If this condition persists, contact NetApp technical support.

**Syslog Message**

Quota scan aborted on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

## quota.scan.cancel

### Severity

NOTICE

### Description

This message occurs after the administrator uses the "volume quota on" command, when the quota scanner is canceled by other activity in the system. An example of such activity would be the "volume quota off" command.

### Corrective Action

(None).

### Syslog Message

Quota scan canceled on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

## quota.scan.complete

### Severity

NOTICE

### Description

This message occurs when the administrator uses the "volume quota on" command and the quota scanner has completed scanning.

### Corrective Action

(None).

### Syslog Message

Quota scan completed on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

## quota.scan.start

### Severity

NOTICE

### Description

This message occurs when the administrator uses the "volume quota on" command and starts the quota scanner.

### Corrective Action

(None).

### Syslog Message

Quota scan started on %s %s%s%s%s.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Additional volume identifier, used to uniquely identify the volume if the volume name alone cannot.

## quota.softlimit events

### quota.softlimit.exceeded

### Severity

NOTICE

### Description

This message occurs when a user exceeds the soft quota limit (file and/or block limit).

### Corrective Action

(None).

### Syslog Message

%s %s for %s on volume %s%s%s%s

### Parameters

**ltype** (STRING): Type of soft limit (file or block).

**verb** (STRING): Error condition (exceeded).

**whom** (STRING): Entity for whom this event is generated.

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**vfname** (STRING): vFiler(tm) name.

## quota.softlimit.normal

### Severity

NOTICE

### Description

This message occurs when a user's soft quota limit (file and/or block limit) returns to normal.

### Corrective Action

(None).

### Syslog Message

%s %s for %s on volume %s%s%s%s

### Parameters

**itype** (STRING): Type of soft limit (file or block).

**verb** (STRING): Status "normal" to indicate that the previously violated condition is now fixed.

**whom** (STRING): Entity for whom this event is generated.

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**vfname** (STRING): vFiler(tm) name.

## quota.upgrade events

### quota.upgrade

### Severity

INFORMATIONAL

### Description

This message occurs when a quota upgrade on the volume is started, ended or aborted.

### Corrective Action

(None).

### Syslog Message

Quota upgrade %s for volume %s%s%s.

### Parameters

**status** (STRING): Indicates that the upgrade has started or is over for this vFiler.

**volname** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which volume name alone is insufficient.

# raid events

## raid.aggr events

### raid.aggr.autoGrow.abort

#### Severity

NOTICE

#### Description

This message occurs when the usable size of an aggregate does not expand after a corresponding increase in the size of the underlying storage.

#### Corrective Action

Check the sizes of all disks in the affected aggregate, and make sure that the backing volumes have increased to the same size. Run the "storage disk show -virtual-machine-disk-info" command to get the name of the backing volume for a disk, under the "VM Disk File" column in the command output.

#### Syslog Message

The capacity of aggregate %s couldn't be expanded. The parameters to increase size are raidgroup\_count: %d, min\_new\_size: %llu, max\_new\_size: %llu. Reason for failure: %s.

#### Parameters

**aggr\_name** (STRING): Name of the aggregate that did not expand in size.  
**rg\_count** (INT): Number of RAID groups that have disks that have increased in size.  
**max\_new\_size** (LONGINT): Maximum possible new usable size across the RAID group  
**min\_new\_size** (LONGINT): Minimum possible new usable size across the RAID group  
**reason** (STRING): Short string describing why the aggregate capacity was not changed.

### raid.aggr.autoGrow.success

#### Severity

NOTICE

#### Description

This message occurs when the usable size of an aggregate is able to expand after a corresponding increase in the size of the underlying storage.

#### Corrective Action

(None).

#### Syslog Message

The capacity of aggregate %s has increased from %llu blocks to %llu blocks.

#### Parameters

**aggr\_name** (STRING): Name of the aggregate that has increased in size.  
**old\_size** (LONGINT): Previous size, in blocks, of the aggregate that has increased in size.  
**new\_size** (LONGINT): New usable size, in blocks, of the aggregate.

## raid.aggr.lock.conflict

### Severity

INFORMATIONAL

### Description

This message occurs when a configuration operation on an aggregate (for example, adding disks, relocating aggregate disks, etc.), fails to get exclusive lock access because of a conflict with an existing configuration operation that is already holding the exclusive lock on the aggregate.

### Corrective Action

Retry the failed operation after the current operation is complete.

### Syslog Message

Exclusive access on the aggregate %s by operation %s failed because of a conflict with current operation %s. Retry after a few minutes.

### Parameters

**aggregate** (STRING): Name of the aggregate.

**attempted\_op\_name** (STRING): Name of the operation that tried to get exclusive access.

**current\_op\_name** (STRING): Name of the operation that is currently holding the exclusive access.

## raid.aggregate events

### raid.aggregate.relocate

#### Severity

INFORMATIONAL

#### Description

This message occurs when the source node prepares an aggregate for relocation to the specified destination node.

#### Corrective Action

(None).

#### Syslog Message

Aggregate %s%s: Prepared for 'relocation' operation.

#### Parameters

**owner** (STRING): String indicating the owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## raid.aggrvote events

### raid.aggrvote.pullNotOk

#### Severity

INFORMATIONAL

## Description

This message occurs when the mirror vote pull from the replicated database (RDB) to the mailbox fails. SyncMirror® aggregates with only one plex will be kept offline. The mirror vote pull is retried automatically.

## Corrective Action

(None).

## Syslog Message

RAID: mirror vote pull from the RDB failed. SyncMirror aggregates with only one plex will be kept offline.

## Parameters

(None).

## raid.aggrvote.updateNotOk

### Severity

ERROR

### Description

This message occurs when the mirror vote update to a replicated database (RDB) fails.

### Corrective Action

The update is retried automatically. If the condition persists, contact NetApp technical support.

### Syslog Message

RAID: mirror vote update to the RDB failed with error %s for aggregate UUID %s (num of plexes %d; ids: %d, %d).

### Parameters

**reason** (STRING): Reason for the failure.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**num\_plexes** (INT): Number of plexes recorded in the mirror vote.

**plex0\_id** (INT): ID of plex 0 in the mirror vote.

**plex1\_id** (INT): ID of plex 1 in the mirror vote.

## raid.aggrvote.updateOk

### Severity

NOTICE

### Description

This message occurs when the mirror vote update to a replicated database (RDB) succeeds.

### Corrective Action

(None).

### Syslog Message

RAID: mirror vote update to the RDB succeeded for aggregate UUID %s (num of plexes %d; ids: %d, %d), action %s.

## Parameters

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.  
**num\_plexes** (INT): Number of plexes recorded in the mirror vote.  
**plex0\_id** (INT): ID of plex 0 in the mirror vote.  
**plex1\_id** (INT): ID of plex 1 in the mirror vote.  
**action\_type** (STRING): Action performed on the RDB table entry. "CREATE" or "MODIFY".

# raid.assim events

## raid.assim.add.offline.disk

### Severity

NOTICE

### Description

This event is generated when adding an offline disk into a raidtree.

### Corrective Action

(None).

### Syslog Message

Adding offline %s to plex\_id:%d

## Parameters

**disk\_info** (STRING): Formatted information of the offline disk.  
**plexId** (INT): The identifier for the plex.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.assim.build.errorDisk

### Severity

NOTICE

### Description

This message occurs when RAID cannot build a RAID tree for a specific plex from a specific disk. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).



## Syslog Message

Error building RAID tree %s, plex %d, from disk %s.

## Parameters

**tree** (STRING): Volume type.  
**plexId** (INT): Identifier for the plex.  
**disk** (STRING): Name of the disk.

## raid.assim.cfo.aggr.offline

### Severity

INFORMATIONAL

### Description

This event occurs when a CFO aggregate which is originally homed to the disaster recovery (DR) partner or DR auxiliary partner is taken offline during assimilation after switchover in a MetroCluster(tm) configuration. Switched over CFO aggregates are kept offline.

### Corrective Action

(None).

## Syslog Message

Aggregate %s%s originally homed to %u is taken offline. Reason: %s.

## Parameters

**owner** (STRING): Owner of affected aggregate.  
**aggregate** (STRING): The name of the aggregate.  
**dr\_home\_owner\_id** (INT): NVRAM system ID of the aggregate's disaster recovery (DR) home owner.  
**reason** (STRING): The reason for taking the aggregate offline.

## raid.assim.cls.mismatchGens

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® excludes a disk from assimilation into this RAID tree because of a configuration error detected by the consistent label set (CLS) algorithm: the disk transaction generation contains an internal inconsistency. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

Data ONTAP takes appropriate recovery actions but if the aggregate is still shown as 'failed' when the command 'sysconfig -r' or 'aggr status -r' is run, then contact NetApp technical support for information about how to resolve the issue.

## Syslog Message

Orphaning disk %s, because of generation count mismatch (%d/%d).

## Parameters

**disk** (STRING): Name of the disk.

**I1gen** (INT): Generation count of the disk's L1 transaction.

**I2gen** (INT): Generation count of the disk's L2 transaction.

## **raid.assim.cls.noStableCIs**

### **Severity**

NOTICE

### **Description**

This message occurs when the consistent label set (CLS) algorithm is restarted because Data ONTAP® could not determine a consistent CLS for all of the RAID aggregates. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

Unable to find a stable CLS for all RAID trees. Resorting disks.

### **Parameters**

(None).

## **raid.assim.cls.noStableCIs2**

### **Severity**

NOTICE

### **Description**

This message occurs when the consistent label set (CLS) algorithm is restarted because Data ONTAP® could not determine a consistent CLS for all of the RAID aggregates with Label 2. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

Unable to find a stable CLS for all RAID trees with Label 2. Resorting disks.

### **Parameters**

(None).

## **raid.assim.cls.retryCIs**

### **Severity**

NOTICE

### **Description**

This message occurs when the consistent label set (CLS) algorithm is restarted due to a disk that appears with an unreasonable timestamp (appears to be a future time). Data ONTAP® takes appropriate recovery actions, as described in additional logged events

**Corrective Action**

(None).

**Syslog Message**

Suspect unreasonable timestamp in future for disk in plex %s/%d. Restarting CLS algorithm.

**Parameters**

**plex** (STRING): Plex name.

**pid** (INT): Plex ID.

**raid.assim.cls.tooManyPlexes****Severity**

NOTICE

**Description**

This message occurs when the number of plexes detected exceeds the maximum limit allowed, during consistent label set (CLS) computation. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

Raidtree %s has more than %d plexes (plex IDs %d, %d, %d)

**Parameters**

**vol** (STRING): Name of the aggregate.

**maxPlexes** (INT): Maximum number of plexes allowed per tree.

**plex1** (INT): Plex #1.

**plex2** (INT): Plex #2.

**plex3** (INT): Plex #3.

**raid.assim.cls.tooRecent****Severity**

NOTICE

**Description**

This message occurs when Data ONTAP® excludes a disk from assimilation into this RAID tree because of a configuration error detected by the consistent label set (CLS) algorithm: the disk transaction identifier indicates that it is more recent than the base transaction ID of the CLS. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

Data ONTAP takes appropriate recovery actions but if the aggregate is still shown as 'failed' when the command 'sysconfig -r' or 'aggr status -r' is run, then contact NetApp technical support for information about how to resolve the issue.

**Syslog Message**

Orphaning disk %s, because its generation (%d/%d) is more recent than %d/%d.

## Parameters

**disk** (STRING): Name of the disk.  
**t1gen** (INT): Generation count of the disk's transaction.  
**t1time** (INT): Timestamp of the disk's transaction.  
**vgen** (INT): Generation count of the aggregate.  
**vtime** (INT): Timestamp of the aggregate transaction.

## raid.assim.cxt.degradedDirty

### Severity

INFORMATIONAL

### Description

This message occurs when RAID detects that a context parity protected aggregate contains a degraded RAID group and has dirty parity.

### Corrective Action

(None).

### Syslog Message

%s "%s%s" is degraded and has dirty parity.

## Parameters

**vol\_type** (STRING): Volume type.  
**host** (STRING): 'Local' or 'Partner' where the aggregate is located.  
**vol** (STRING): Name of the aggregate.

## raid.assim.disk.badlabelversion

### Severity

ALERT

### Description

This message occurs when RAID assimilation discovers that a disk has an invalid raid label version. Data ONTAP® marks the disk as broken.

### Corrective Action

Contact NetApp technical support. For information about correcting the problem, search for the "raid.assim.disk.badlabelversion" keyword on the Knowledgebase of the NetApp Support Site.

### Syslog Message

%s has %s RAID label with version (%d), which is not within the currently supported range (%d - %d).

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM and carrier serial number.  
**gen** (STRING): Non-empty string for the old format label.  
**version** (INT): Label version number.  
**versionLow** (INT): Lower range of the valid label version.  
**versionHigh** (INT): Upper range of the valid label version.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.assim.disk.broken**

### **Severity**

ERROR

### **Description**

This message occurs when we discover that a broken disk has been discovered during assimilation.

### **Corrective Action**

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### **Syslog Message**

Broken %s detected during assimilation.

### **Parameters**

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.assim.disk.brokenPreAssim**

### **Severity**

ERROR

### **Description**

This message occurs when we discover that a broken disk has been discovered prior to assimilation.

### **Corrective Action**

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be

issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### Syslog Message

Broken %s detected prior to assimilation.

### Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.assim.disk.md5error

### Severity

NOTICE

### Description

This event occurs when we detect that a disk label has an MD5 signature that is different from the MD5 signature of the disk that it's on. Usually this is due to a disk serial number change.

### Corrective Action

(None).

### Syslog Message

%s%s%s, diskobj\_verify: MD5 signature of disk %s (S/N %s) does not match with its label signature. This has been corrected.

### Parameters

**container** (STRING): The type of container (aggregate, disk pool).  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk** (STRING): The name of the disk  
**serial** (STRING): The serial number of the disk  
**oldmd5** (STRING): The old MD5 checksum of the disk  
**newmd5** (STRING): The new MD5 checksum of the disk

## raid.assim.disk.nolabels

### Severity

EMERGENCY

## Description

This event occurs when we detect that a disk has no valid labels. Disks in NetApp systems should always have valid labels; therefore, a disk without labels implies that there is corruption somewhere in the data path to the disk. (This is the most common cause, often because of loose cables or a newly added shelf.) It is possible, although much less likely, that the disk itself is experiencing silent data corruption. The affected disk is taken out of service for the reason "bad label."

## Corrective Action

If the disk was newly added and really should be treated as a spare, the "disk unfail" command will mark the disk as a spare. Otherwise, if the disk was previously working normally, check the connectivity to the disk. Are all the cables secured? Sometimes newly added shelves or loose cables can cause electrical noise on the bus that corrupts data transmission. If the connectivity to the disk is good and there is filesystem data on the disk, contact NetApp technical support for help.

## Syslog Message

%s has no valid labels. It will be taken out of service to prevent possible data loss.

## Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.assim.disk.region.hole

### Severity

ALERT

### Description

This message occurs when RAID assimilation detects that a disk has a hole between the two disk regions. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

Contact NetApp technical support. For information about correcting the problem, search for the "raid.assim.disk.region.hole" keyword on the Knowledgebase of the NetApp Support Site.

### Syslog Message

Disk %s (S/N %s) has a hole in its allocated space between region %d (type %s,start %llu, size %llu) and region %d (type %s, start %llu, size %llu).

### Parameters

**disk** (STRING): Name of the disk.

**diskSerialno** (STRING): Serial number of the disk.

**regionId1** (INT): First disk region identifier.

**region1** (STRING): First disk region name.  
**start1** (LONGINT): Start of the first disk region.  
**size1** (LONGINT): Size of the first disk region.  
**regionId2** (INT): Second disk region identifier.  
**region2** (STRING): Second disk region name.  
**start2** (LONGINT): Start of the second disk region.  
**size2** (LONGINT): Size of the second disk region.

## **raid.assim.disk.region.overlap**

### **Severity**

ERROR

### **Description**

This message occurs when RAID assimilation detects that a disk has two regions overlapping. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

Contact NetApp technical support. For information about correcting the problem, search for the "raid.assim.disk.region.overlap" keyword on the Knowledgebase of the NetApp Support Site.

### **Syslog Message**

On disk %s (S/N %s), region %d (type %s, start %llu, size %llu) overlaps with region %d (type %s, start %llu, size %llu).

### **Parameters**

**disk** (STRING): Name of the disk.  
**diskSerialNo** (STRING): Serial number of the disk.  
**regionId1** (INT): First disk region identifier.  
**region1** (STRING): First disk region name.  
**start1** (LONGINT): Start of the first disk region.  
**size1** (LONGINT): Size of the first disk region.  
**regionId2** (INT): Second disk region identifier.  
**region2** (STRING): Second disk region name.  
**start2** (LONGINT): Start of the second disk region.  
**size2** (LONGINT): Size of the second disk region.

## **raid.assim.disk.rightsizeChange**

### **Severity**

INFORMATIONAL

### **Description**

This event indicates that the rightsize stored in this disk's Table Of Contents(TOC) is different from the rightsize currently being reported by the disk.

### **Corrective Action**

(Call support).

### **Syslog Message**

(None).



## Parameters

**diskType** (STRING): Indicates if this a spare or filesystem disk.  
**disk\_info** (STRING): Formatted information of the disk  
**toc\_rightsize** (LONGINT): The disk rightsize according to the disk TOC.  
**changeType** (STRING): Type of change: growth or shrinkage.  
**storage\_rightsize** (LONGINT): The disk rightsize according to the storage layer.  
**systemAction** (STRING): The action being taken by the system to deal with this event.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm)(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.assim.disk.spare

### Severity

NOTICE

### Description

This event is generated when we spare an error disk.

### Corrective Action

(None).

### Syslog Message

Sparing %s, because %s

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**reason** (STRING): (None).  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.assim.fatal

### Severity

EMERGENCY

## Description

This message occurs when assimilation encounters a fatal error, such as running out of memory.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Assimilation failed: %s

## Parameters

**reason** (STRING): Reason assimilation failed.

## raid.assim.fatal.upgrade

### Severity

EMERGENCY

## Description

This message occurs when RAID assimilation encounters a fatal error during an unsuccessful upgrade to a newer version of Data ONTAP®.

## Corrective Action

Contact NetApp technical support. For information about correcting the problem, search for the "raid.assim.fatal" keyword on the Knowledgebase of the NetApp Support Site. You can also boot with the previous version of Data ONTAP.

## Syslog Message

Assimilation failed: This system appears to be upgrading from a previous version of Data ONTAP. Errors were found while upgrading the aggregates. The upgrade will abort now.

## Parameters

(None).

## raid.assim.label.makeForeignVol

### Severity

ERROR

## Description

This message occurs when ownership cannot be claimed over a previously non-native aggregate.

## Corrective Action

Use the 'storage aggregate online' command to bring the aggregate online.

## Syslog Message

%s %s could not be turned into a %snative aggregate; it will appear to be foreign.

## Parameters

**vol\_type** (STRING): Volume type. Always aggregate.

**vol** (STRING): Name of the aggregate.

**host** (STRING): Empty string for local aggregates, or "partner" for partner aggregates.

**vol\_type2** (STRING): No longer used.

## **raid.assim.label.makeNativeVol**

### **Severity**

NOTICE

### **Description**

This event is generated when we claim ownership over a previously non-native aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

%s %s turned into a %snative %s.

### **Parameters**

**vol\_type** (STRING): Volume type.

**vol** (STRING): Name of the aggregate.

**host** (STRING): Empty string for local aggregates, or "partner " for partner aggregates.

**vol\_type2** (STRING): Volume type.

## **raid.assim.label.noNativeVols**

### **Severity**

NOTICE

### **Description**

This event is generated when we complete assimilation and discover that none of the volumes or aggregates match the local system id. This can occur if the NVRAM card has been swapped. The system will attempt to select a set of volumes and aggregates to make native.

### **Corrective Action**

(None).

### **Syslog Message**

No native volumes or aggregates detected, assuming NVRAM card swap. %d NVRAM system id(s) (%u, %u, %u...) found among all volumes. NVRAM system id %u selected. NVRAM card swap detected by %s.

### **Parameters**

**nsets** (INT): Number of different serial numbers found in non-native volumes.

**nvrasm\_sysid1** (INT): First NVRAM system id to select from.

**nvrasm\_sysid2** (INT): Second NVRAM system id to select from.

**nvrasm\_sysid3** (INT): Third NVRAM system id to select from.

**selected\_nvrasm\_sysid** (INT): Selected NVRAM system id of volumes before swap.

**who** (STRING): Module detected NVRAM card swap.

## **raid.assim.label.upgrade**

### **Severity**

INFORMATIONAL

**Description**

This event is generated when we upgrade RAID labels to the new 6.2-and-later format.

**Corrective Action**

(None).

**Syslog Message**

Upgrading RAID labels.

**Parameters**

(None).

**raid.assim.label.upgrade.corruptSize****Severity**

ALERT

**Description**

This message occurs when an old RAID label is being upgraded and the label appears to be corrupted due to the parity type and the sizes[0] entry (describing the size of the parity disk) conflicting. This deals with burt 79966.

**Corrective Action**

Wait for the event message raid.disk.unload.done or raid.carrier.remove to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, raid.disk.unload.done will be issued immediately. Otherwise, wait for raid.carrier.remove which indicates the carrier has been fully prepared for removal. Contact NetApp technical support.

**Syslog Message**

Disk %s has a corrupt old label: sizes[%d] is %u.

**Parameters**

**disk** (STRING): Disk with the corrupted label.  
**rgdn** (INT): Disk number given by the RAID Group.  
**size** (INT): Size of the disk.

**raid.assim.label.upgrade.parityMismatch****Severity**

ALERT

**Description**

This message occurs when an old RAID label is being upgraded and the label appears to be corrupted due to the parity type and the sizes[0] entry (describing the size of the parity disk) conflicting. This deals with burt 79966.

**Corrective Action**

Wait for the event message raid.disk.unload.done or raid.carrier.remove to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, raid.disk.unload.done will be issued immediately. Otherwise, wait for raid.carrier.remove which indicates the carrier has been fully prepared for removal. Contact NetApp technical support.

## Syslog Message

Disk %s has a corrupt old label: parity type mismatch.

## Parameters

**disk** (STRING): Disk with the corrupted label.

## raid.assim.mark.error.disk.broken

## Severity

NOTICE

## Description

This message occurs when RAID moves an invalid disk into a broken disk pool. Invalid disks can be ones with bad labels or with an unsupported label version.

## Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

## Syslog Message

Marking disk %s, which failed assimilation in %s %s, as broken, because %s.

## Parameters

**disk** (STRING): Name of the disk.

**vol\_type** (STRING): Volume type.

**tree** (STRING): Name of the aggregate.

**reason** (STRING): Reason for marking the disk broken.

## raid.assim.rg.dupChildId

## Severity

ALERT

## Description

This event occurs when we detect that a RAID group object has an inconsistent list of disk objects in its label configuration record.

## Corrective Action

(Call support).

## Syslog Message

%s %s%s, rgobj\_verify: RAID object %d has duplicate child objects (%s and %s) with ID %d.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**rgId** (INT): The RAID group identifier

**diskName1** (STRING): One disk with this identifier

**diskName2** (STRING): Another disk with this identifier

**childId** (INT): The child identifier

## **raid.assim.rg.offlineVerifyFailed**

### **Severity**

NOTICE

### **Description**

This message occurs when an offline disk within a RAID group fails verification checks. For example, ownership information differs with the parent aggregate, or the parent tree map specifies this disk is unmapped. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

%s "%s%s", Offline child object (ID=%d) in RAID object %d failed verification

### **Parameters**

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**rgId** (INT): RAID group identifier.  
**childId** (INT): Child identifier.

## **raid.assim.tree.degradedDirty**

### **Severity**

EMERGENCY

### **Description**

This message occurs when an aggregate contains a degraded RAID group and has dirty parity. Wafiron must be run before the aggregate can be brought online.

### **Corrective Action**

Run wafiron for the aggregate by using the command 'aggr wafiron start'.

### **Syslog Message**

%s "%s%s" is degraded and has dirty parity. You must run wafiron on this aggregate before it can be brought online.

### **Parameters**

**vol\_type** (STRING): Volume type.  
**host** (STRING): 'Local' or 'Partner' where the aggregate is located.  
**vol** (STRING): Name of the aggregate.  
**vol\_type2** (STRING): Volume type.

## **raid.assim.tree.dupld**

### **Severity**

ALERT

## Description

This message occurs when RAID assimilation detects that a duplicate volume identifier exists. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

## Corrective Action

Contact NetApp technical support. For information about correcting the problem, search for the "raid.assim.tree.dupId" keyword on the Knowledgebase of the NetApp Support Site.

## Syslog Message

%s %s%s and %s %s%s have the same RAID tree ID. RAID assimilation will be aborted.

## Parameters

**vol\_type1** (STRING): Volume type.

**owner1** (STRING): Owner of the first aggregate.

**vol1** (STRING): Name of the first aggregate.

**vol\_type2** (STRING): Volume type.

**owner2** (STRING): Owner of the second aggregate.

**vol2** (STRING): Name of the second aggregate.

## raid.assim.tree.dupName

### Severity

NOTICE

### Description

This message occurs when the system detects a duplicate name for a volume or an aggregate. In this case, the second volume or aggregate with that name is renamed.

### Corrective Action

Use the 'vol rename' command or 'aggr rename' command to resolve the conflict permanently.

### Syslog Message

Duplicate %s names found, an instance of %s%s is being renamed to %s%s.

### Parameters

**vol\_type** (STRING): Whether the duplicate name is for a volume or an aggregate.

**owner1** (STRING): String indicating the owner of the volume or aggregate.

**vol** (STRING): Original name of the volume or aggregate.

**owner2** (STRING): String indicating the owner of the volume or aggregate.

**newName** (STRING): New name of the volume or aggregate.

## raid.assim.tree.foreign

### Severity

ERROR

### Description

This message occurs when a RAID tree object contains an ownership identifier that is different from the local system. In this case, the aggregate is declared "foreign" and is marked offline.

### Corrective Action

Use the 'storage aggregate online' command to bring the aggregate online.

### Syslog Message

raidtree\_verify: %s %s is a foreign aggregate and is being taken offline. Use the 'storage aggregate online' command to bring it online.

### Parameters

**vol\_type** (STRING): Volume type. Always aggregate.

**vol** (STRING): Name of the aggregate.

**vol\_type2** (STRING): No longer used.

**vol\_type3** (STRING): No longer used.

## raid.assim.tree.multipleRootVols

### Severity

EMERGENCY

### Description

This message occurs when the system detects an assimilation error indicating that two aggregates claim to be the root volume.

### Corrective Action

Go into Maintenance mode and use the 'aggr options' command with the 'root' option to pick a root aggregate.

### Syslog Message

Volume %s and volume %s both claim to be the %sroot volume.

### Parameters

**vol1** (STRING): Name of the first aggregate.

**vol2** (STRING): Name of the second aggregate.

**host** (STRING): Empty string for local aggregate, or "partner " for partner aggregate.

## raid.assim.tree.noRootVol

### Severity

ERROR

### Description

This message occurs when the system detects an assimilation error indicating that no root aggregate is found.

### Corrective Action

Go into Maintenance mode and use the 'aggr options' command with the 'root' option to pick an aggregate to be the root volume.

### Syslog Message

No usable %sroot volume was found!



## Parameters

**host** (STRING): Empty string for local aggregates, or "partner " for partner aggregates.

## raid.assim.tree.tooManyChild

### Severity

ALERT

### Description

This event occurs when we detect that a RAID tree object has more than one root volume objects in its label configuration record.

### Corrective Action

(Call support).

### Syslog Message

%s %s%s,raidtree\_verify: Raidtree has multiple root objects.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.assim.upgrade.aggr.fail

### Severity

EMERGENCY

### Description

This message occurs when assimilation is attempting to upgrade the RAID labels for all aggregates (usually because of a kernel upgrade) and a aggregate can not be successfully upgraded. The RAID label upgrade is aborted. Boot with the previous release to fix the problem with the aggregate (which should be noted in the boot-time messages).

### Corrective Action

Boot with the previous version of Data ONTAP to fix the problem with the aggregate (which should be noted in the boot-time messages).

### Syslog Message

RAID label upgrade of %s %s%s failed.

### Parameters

**voltype** (STRING): Type of RAID volume (volume or aggregate)

**owner** (STRING): Whether this is a partner or local volume

**volume** (STRING): Name of the volume that had a problem.

## raid.autopart events

## raid.autoPart.disabled

### Severity

ALERT

### Description

This message occurs when auto-partitioning for root aggregate creation is disabled on a system. When auto-partitioning is disabled, the root aggregate is created on whole disks. Systems with a root aggregate on partitioned disks can devote more space to non-root aggregates than systems with a root aggregate on whole disks. If this message occurs after root aggregate creation, then further partitioning of the disks is disabled.

### Corrective Action

If the system is already initialized with root aggregate on partitioned disks, ONTAP will try to re-enable auto-partitioning after the system boots up and the root aggregate is online. The event `raid.autoPart.enabled` will be logged when auto-partitioning is re-enabled. To create a root aggregate on partitioned disks: 1. Halt the controller. 2. Replace any failed internal disks (see the "reason" field in the EMS message). 3. Reboot the controller. 4. Re-initialize the system. If this message is seen after the system initialization and the root aggregate is already created on partitioned disks, then the system will continue to function, as follows: 1. Expansion of existing aggregates created from partitioned disks may not work. 2. If a disk fails in a RAID group created out of partitioned disks, then a spare will fail to be partitioned and disk reconstruction will not start. 3. If a disk is marked for predictive failure in a RAID group created out of partitioned disks, then a spare will fail to be partitioned and the sick disk copy will not start. 4. If disks are available, then manually partition the disks and keep them as spares. 5. If you were previously directed to disable ADP by NetApp technical support, reach out to NetApp technical support to re-enable ADP.

### Syslog Message

Disk auto-partitioning is disabled on this system: %s.

### Parameters

**reason** (STRING): Reason that auto-partitioning was disabled.

## raid.autoPart.done

### Severity

NOTICE

### Description

This message indicates that auto-partitioning has completed.

### Corrective Action

(None).

### Syslog Message

Successfully auto-partitioned %d of %d disks.

### Parameters

**n\_partitioned** (INT): Number of disks successfully auto-partitioned.

**n\_total** (INT): Number of disks available for auto-partitioning.

## raid.autoPart.enabled

### Severity

NOTICE

### Description

This message occurs when auto-partitioning for root volume creation is enabled on the system.

### Corrective Action

None.

### Syslog Message

Auto-partitioning is enabled on the system %s for partition method %s with %llu root size blocks.

### Parameters

**node** (STRING): Node on which auto-partitioning is enabled.

**partition\_method** (STRING): Partition method which is enabled.

**size\_blocks** (LONGINT): Thin partition size in blocks for root.

## raid.autoPart.paramsModified

### Severity

NOTICE

### Description

This message occurs when the "root-configuration" environment variable, which controls auto-partitioning for the initial root volume, contains one or more invalid parameters. When invalid parameters are discovered, the system reverts to using default parameters instead of the invalid parameters.

### Corrective Action

Use the suggested "root-configuration" setting instead of the existing "root-configuration" setting.

### Syslog Message

Invalid "root-configuration" settings: "%s". The system will auto-partition as if "root-configuration" had been set to: "%s".

### Parameters

**original\_setting** (STRING): Original environment variable.

**modified\_setting** (STRING): Modified environment variable.

## raid.autoPart.start

### Severity

NOTICE

### Description

This message indicates that auto-partitioning has started.

### Corrective Action

(None).

## Syslog Message

System has started auto-partitioning %d disks.

## Parameters

**n\_total** (INT): Number of disks available for auto-partitioning.

# raid.capacity.events

## raid.capacity.license.alloc

### Severity

NOTICE

### Description

This message occurs when the system creates a new aggregate, or registers a new capacity for an existing aggregate.

### Corrective Action

(None).

## Syslog Message

Registered licensed capacity (%llu%s) for %s aggregate %s UUID %s - operation: %s, simulated: %d.

## Parameters

**capacity** (LONGINT): Capacity registered for the aggregate against the license manager.

**capacity\_unit** (STRING): Capacity unit (MB, GB, TB) for registered capacity.

**mirrored** (STRING): "mirrored" if the aggregate is mirrored; "unmirrored" otherwise.

**aggregate\_name** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**operation** (STRING): Aggregate operation triggering the license capacity update.

**simulated** (INT): 1 if the operation is simulated; 0 if the operation is not simulated.

**root** (INT): 1 if the aggregate is a root aggregate; 0 if it is a data aggregate.

**status** (INT): Synchronous status returned by the license capacity alloc call.

**callback** (INT): Indicates whether callback requested

**sequence** (INT): Sequence number to prevent suppression of duplicate messages.

## raid.capacity.license.free

### Severity

NOTICE

### Description

This message occurs when the system destroys an aggregate, and thus frees its registered license capacity.

### Corrective Action

(None).

## Syslog Message

Freed licensed capacity of aggregate %s UUID %s operation: %s.

## Parameters

**aggregate\_name** (STRING): Name of the aggregate.  
**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.  
**operation** (STRING): Aggregate operation triggering the license capacity free.  
**status** (INT): Synchronous status returned by the license capacity free call.  
**sequence** (INT): Sequence number to prevent suppression of duplicate messages.

## raid.capacity.license.status

### Severity

NOTICE

### Description

This message occurs when the system attempts to update the licensed capacity of the system due to an aggregate operation.

### Corrective Action

(None).

### Syslog Message

Status of licensing request %s, aggregate %s, licensed capacity %llu%s, capacity remaining %llu%s, additional capacity required %llu%s.

## Parameters

**status\_string** (STRING): Status of the licensing request.  
**aggregate\_name** (STRING): Name of the aggregate.  
**licensed\_size** (LONGINT): Licensed capacity of the aggregate.  
**licensed\_size\_unit** (STRING): Capacity unit (MB, GB, TB) for licensed capacity.  
**capacity\_left** (LONGINT): Capacity remaining in the license.  
**capacity\_left\_unit** (STRING): Capacity unit (MB, GB, TB) for capacity remaining.  
**capacity\_required** (LONGINT): Additional licensed capacity required for this operation.  
**capacity\_required\_unit** (STRING): Capacity unit (MB, GB, TB) for additional capacity required.  
**status\_code** (INT): Specific status code returned by the licensing request.  
**sequence** (INT): Sequence number to prevent suppression of duplicate messages.

## raid.carrier events

### raid.carrier.diskEvacuated

### Severity

NOTICE

### Description

This message occurs when the evacuation operation is complete on a disk in a multidisk carrier that contains at least one failed disk. When a disk in a multidisk carrier fails, Data ONTAP® initiates evacuation of the other disks in the carrier so that you can replace the carrier after the evacuation is done. The disk's amber Attention LED is lit continuously after the carrier evacuation is complete.

### Corrective Action

(None).

## Syslog Message

Disk evacuation operation is complete on %s.

## Parameters

**disk\_info** (STRING): Formatted information of the disk being evacuated.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of the disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.carrier.doNotRemove

### Severity

NOTICE

### Description

This message occurs when a single disk fails while it is in a carrier that contains more than one disk. This disk carrier should not be removed because another disk might be in use. If another disk is in use, Data ONTAP® performs a disk copy operation on the other disk so that the carrier can be replaced after the disk copy is completed.

### Corrective Action

Do not remove the disk carrier with the failed disk until the disk copy operation is completed. You can remove the carrier after the 'raid.carrier.remove' event is issued.

## Syslog Message

%s in a multi-disk carrier failed. DO NOT REMOVE the disk carrier with the failed disk until the 'raid.carrier.remove' event is issued.

## Parameters

**disk\_info** (STRING): Formatted information of the disk object.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of the disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.carrier.remove

### Severity

ALERT

### Description

This message occurs when the last disk in a carrier that contains more than one disk has failed.

### Corrective Action

Replace the carrier containing the failed disks.

### Syslog Message

%s failed after all other disks have failed in the carrier. Replace the carrier.

### Parameters

**disk\_info** (STRING): Information about the last failed disk, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): String representing the model type of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is "Local".

## raid.cksum events

### raid.cksum.bad.file.block

### Severity

ERROR

### Description

This event is issued when WAFL accesses a block that was zeroed because it could not be recovered correctly.

### Corrective Action

Please restore the corrupted file from snapshot or from the backup media.

### Syslog Message

Reading bad block on %s %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

### Parameters

**vol\_type** (STRING): Whether the block was being accessed for a flexible volume or an aggregate.

**owner** (STRING): Owner of the affected volume or aggregate.

**vol** (STRING): The name of the volume or aggregate

**disk\_info** (STRING): Formatted information of the disk

**blockNum** (LONGINT): The block number  
**ino\_type** (STRING): The inode type  
**fileid** (INT): The file identifier  
**snapid** (INT): The snapshot identifier  
**bno** (LONGINT): The block number in the file  
**level** (INT): The level  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.buf.read.error.embedded

### Severity

NOTICE

### Description

This message occurs when the system detects an embedded checksum error. This can occur when a zeroing disk is added to the RAID group for reconstruction, and the zeroing is never completed.

### Corrective Action

(None).

### Syslog Message

Embedded checksum error on %s %s%s, %s, phys checksum block %llu

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Formatted information about the disk.  
**blockNum** (LONGINT): Disk block number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model name of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.



## raid.cksum.bypass

### Severity

NOTICE

### Description

This message occurs when a read error has occurred while media error bypass is enabled.

### Corrective Action

(None).

### Syslog Message

raid\_cksum\_compare: media error bypass enabled; Ignoring read error - %s %s%s, %s disk block %llu

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Formatted information about the disk.

**blockNum** (LONGINT): Block number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.bypass.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch while media error bypass is enabled. The expected virtual block number (VBN)/disk block number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. The error is ignored.

### Corrective Action

(None).

### Syslog Message

Block number mismatch: %s %s%s, %s stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu; media error bypass enabled; ignoring read error.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Formatted information about the disk.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model name of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.bypass.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry while media error bypass is enabled. The checksum computed over the checksum entry does not match the embedded checksum, indicating that the checksum entry is corrupted. The error is ignored.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry: %s %s%s, %s disk block %llu; Media error bypass enabled; Ignoring read error.

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Formatted information about the disk.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model name of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.bypass.unknown

### Severity

NOTICE

### Description

This message occurs when a read error has occurred and has been ignored because media error bypass is enabled.

### Corrective Action

(None).

### Syslog Message

raid\_cksum\_compare: media error bypass enabled; Ignoring read error - %s %s%s, on disk block %llu

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**blockNum** (LONGINT): Block number.

## raid.cksum.bypass.unkwn.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch while media error bypass is enabled. The expected virtual block number (VBN)/disk block number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. The error is ignored. Details on which disk the error was encountered on are not Available because the disk transitioned to an unavailable state between the detection of the error and the logging of this message.

### Corrective Action

(None).

### Syslog Message

Block number mismatch: %s %s%s, stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu; Media error bypass enabled; Ignoring read error.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_dbn** (LONGINT): Expected physical disk block number.

**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected volume block number.

## raid.cksum.bypass.unkwn.embd

**Severity**

NOTICE

**Description**

This message occurs when the system detects an invalid checksum entry while media error bypass is enabled. The embedded checksum computed over the checksum entry does not match the stored embedded checksum in the checksum entry, indicating the corruption of the checksum entry. The error is ignored. Details on which disk the error was encountered on are not Available, because the disk transitioned to an unavailable state between the detection of the error and the logging of this message.

**Corrective Action**

(None).

**Syslog Message**

Invalid checksum entry: %s %s%s, on block %llu; Media error bypass enabled; Ignoring read error.

**Parameters****vol\_type** (STRING): Volume type.**owner** (STRING): Owner of the affected aggregate.**vol** (STRING): Name of the aggregate.**blockNum** (LONGINT): Physical disk block number containing the error.**raid.cksum.bypass.unkwn.zero****Severity**

NOTICE

**Description**

This message occurs when the system detects an empty checksum entry while media error bypass is enabled. The checksum entry is zeroed, but the corresponding block is not zeroed. The error is ignored. Details on which disk the error was encountered on are not Available because the disk transitioned to an unavailable state between the detection of the error and the logging of this message.

**Corrective Action**

(None).

**Syslog Message**

Empty checksum entry for non-zeroed block: %s %s%s, on disk block %llu; Media error bypass enabled; Ignoring read error.

**Parameters****vol\_type** (STRING): Volume type.**owner** (STRING): Owner of the affected aggregate.**vol** (STRING): Name of the aggregate.**blockNum** (LONGINT): Physical disk block number containing the error.**raid.cksum.bypass.zero****Severity**

NOTICE

## Description

This message occurs when the system detects an empty checksum entry while media error bypass is enabled. The checksum entry is zeroed, but the corresponding block is not zeroed. The error is ignored.

## Corrective Action

(None).

## Syslog Message

Empty checksum entry for non-zeroed block: %s %s%s, %s disk block %llu; Media error bypass enabled; Ignoring read error.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Formatted information about the disk.

**blockNum** (LONGINT): Physical disk block number containing the error.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.disabled

### Severity

NOTICE

### Description

This event is issued when we disable checksums on an aggregate.

### Corrective Action

(None).

### Syslog Message

Disabled checksums on %s %s%s.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.cksum.embedded.error

**Severity**

INFORMATIONAL

**Description**

This message occurs when the stored checksum of a checksum block does not match the computed checksum of the checksum block. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

Checksum Entry of volume %s, disk %s, disk block %llu is bad: expected checksum %x, actual checksum %x

**Parameters**

**vol** (STRING): Name of the aggregate.  
**disk** (STRING): Name of the disk.  
**blockNum** (LONGINT): Disk block number.  
**expected\_cksum** (INT): Expected checksum of the checksum entry.  
**actual\_cksum** (INT): Actual embedded checksum in the checksum entry.

**raid.cksum.enabled****Severity**

NOTICE

**Description**

This event is issued when we enable checksums on an aggregate.

**Corrective Action**

(None).

**Syslog Message**

Enabled checksums on %s %s%s.

**Parameters**

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

**raid.cksum.err.recFail****Severity**

NOTICE

**Description**

This message occurs when a disk is recommended for failure because we encountered more than allowed number of checksum errors.

## Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

## Syslog Message

Disk recommended for failure for exceeding threshold of %d checksum errors on %s%s.

## Parameters

**threshold** (INT): The threshold for checksum errors  
**owner** (STRING): String indicating the owner of the affected volume.  
**disk\_info** (STRING): The name of the disk.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.fsm.phase1

### Severity

INFORMATIONAL

### Description

This event is generated when phase 1 of a checksum upgrade is complete.

### Corrective Action

(None).

### Syslog Message

Completed RAID checksum upgrade phase 1 (of 2) on %s %s%s; phase 2 will occur during the next disk scrub.

### Parameters

**vol\_type** (STRING): Volume type.  
**host** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.cksum.fsm.phase1.abort

### Severity

INFORMATIONAL

## Description

This event is generated when phase 1 of a checksum upgrade is aborted.

## Corrective Action

(None).

## Syslog Message

RAID checksum upgrade phase 1 (of 2) on %s %s%s %s aborted.

## Parameters

**vol\_type** (STRING): Volume type.

**host** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Formatted information of disk object that has completed its zeroing operation

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.fsm.reinit

### Severity

INFORMATIONAL

### Description

This event is generated when reinitialization of checksums is complete.

### Corrective Action

(None).

### Syslog Message

Completed RAID checksum reinitialization on %s %s%s.

### Parameters

**vol\_type** (STRING): Volume type.

**host** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.cksum.fsm.reinit.abort

### Severity

INFORMATIONAL



## Description

This event is generated when checksum reinitialization is aborted.

## Corrective Action

(None).

## Syslog Message

RAID checksum reinitialization on %s %s%s %s aborted.

## Parameters

**vol\_type** (STRING): Volume type.

**host** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Formatted information of disk object that has completed its zeroing operation

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.fsm.revert

### Severity

INFORMATIONAL

### Description

This event is generated when a checksum revert is complete.

### Corrective Action

(None).

### Syslog Message

Checksum revert of %s %s%s completed; vol copy, aggr copy and snapmirror now allowed.

### Parameters

**vol\_type** (STRING): Volume type.

**host** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.cksum.ignore.bno

### Severity

NOTICE

## Description

This message occurs when the system detects a block number mismatch. The expected virtual block number (VBN)/disk block number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. This error is ignored if one of the following conditions is true: - It was encountered during WAFL® check. - It was encountered during wafiron and in ignore unrecoverable checksum error mode. - The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. - It was encountered in ignore checksum error mode. - It was encountered in boot bypass media error mode.

## Corrective Action

(None).

## Syslog Message

Block number mismatch on %s %s %s, %s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu. Ignoring unrecoverable error.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_dbn** (LONGINT): Expected physical disk block number.

**stored\_vbn** (LONGINT): VBN stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected VBN.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.ignore.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry. The embedded checksum computed over the checksum entry does not match, indicating the corruption of the checksum entry. This error is ignored if one of the following conditions is true: - It was encountered during WAFL® check. - It was encountered during wafiron and in ignore unrecoverable checksum error mode. - The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. - It was encountered in ignore checksum error mode. - It was encountered in boot bypass media error mode.

## Corrective Action

(None).

## Syslog Message

Invalid checksum entry on %s %s %s, %s, disk block %llu. Ignoring unrecoverable error.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical disk block number of the disk containing the error.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.ignore.error

### Severity

NOTICE

### Description

This message occurs when the system ignores a checksum error. This error is ignored if one of the following conditions is true: \* It was encountered during WAFL® check. \* It was encountered during wafiron and in ignore unrecoverable checksum error mode. \* The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. \* It was encountered in ignore checksum error mode. \* It was encountered in boot bypass media error mode.

### Corrective Action

(None).

### Syslog Message

%s on %s %s%s, %s, disk block %llu. Ignoring unrecoverable error.

### Parameters

**error** (STRING): Name of the error.

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Block number.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model name of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.ignore.file.bno

### Severity

ALERT

### Description

This message occurs when the system detects a block number mismatch in a block of a file. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry indicating that the block is read from the wrong location. This error is ignored if one of the following conditions is true: - It was encountered during WAFL® check. - It was encountered during wafiron and in ignore unrecoverable checksum error mode. - The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. - It was encountered in ignore checksum error mode. - It was encountered in boot bypass media error mode.

### Corrective Action

Contact technical support for information about how to resolve the unrecoverable checksum error in the referenced file.

### Syslog Message

Block number mismatch on %s %s %s, %s, stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu, %s inode number %d, snapid %d, file block %llu, level %d. Ignoring unrecoverable error.

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**ino\_type** (STRING): The inode type  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot(tm) copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Level.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.ignore.file.embed

### Severity

ERROR

### Description

This message occurs when the system detects an invalid checksum entry in a block of a file. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. This error is ignored if one of the following conditions is true: - It was encountered during WAFL® check. - It was encountered during wafiron and in ignore unrecoverable checksum error mode. - The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. - It was encountered in ignore checksum error mode. - It was encountered in boot bypass media error mode.

### Corrective Action

Contact technical support for information about how to resolve the unrecoverable checksum error in the referenced file.

### Syslog Message

Invalid checksum entry on %s %s %s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d. Ignoring unrecoverable error.

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**ino\_type** (STRING): The inode type  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot(tm) copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Level.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.ignore.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry. The checksum entry is zeroed, but the corresponding block is not zeroed. This error is ignored if one of the following conditions is true: - It was encountered during WAFL® check. - It was encountered during wafiron and in ignore unrecoverable checksum error mode. - The block cannot be recovered from parity and the volume is ignoring unrecoverable errors. - It was encountered in ignore checksum error mode. - It was encountered in boot bypass media error mode.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s %s %s, %s, disk block %llu. Ignoring unrecoverable error.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical disk block number of the disk containing the error.

**shelf** (STRING): Identifier of the disk shelf where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): String representing the name of the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.nvram.error

### Severity

NOTICE

### Description

This message occurs when we cannot write out the checksum for an aggregate because of uninitialized NVRAM. System panics because of this error and will attempt to perform recovery in the next boot.

### Corrective Action

(None).

## Syslog Message

Cannot write checksum for %s %s%s, %s, disk block %llu: NVRAM is not initialized.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Information about the disk object including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Block number.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): String representing the name of the model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.nvram.uninit

### Severity

NOTICE

### Description

This message occurs when an aggregate checksum cannot be repaired because of uninitialized NVRAM.

### Corrective Action

(None).

## Syslog Message

Cannot repair checksum for %s %s%s, %s, disk block %llu: NVRAM is not initialized.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Block number.  
**shelf** (STRING): Identifier of the disk shelf where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): String representing the name of the model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is

located. For non-MetroCluster configurations, site is 'Local'.

## **raid.cksum.replay.bad.entry**

### **Severity**

ERROR

### **Description**

This event is issued at the end of checksum replay when bad checksum entries are detected in NVRAM during that replay (indicating corrupted NVRAM). Reparity will start automatically and update checksum blocks involved in recent writes.

### **Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

### **Syslog Message**

Skipped %d bad checksum entries in NVRAM.

### **Parameters**

**nentries** (INT): The number of bad entries

## **raid.cksum.replay.disk.missing**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when we encounter an NVRAM checksum entry for a missing disk during checksum replay. The disk will be reconstructed if an appropriate spare is available.

### **Corrective Action**

(None).

### **Syslog Message**

Skipping checksum block %llu because disk with index %d is missing in %s %s%s, plex %s, RAID group %d.

### **Parameters**

**cbn** (LONGINT): Checksum block number  
**diskid** (INT): Index of the disk in the RAID group.  
**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**plex** (STRING): The name of the plex.  
**rgid** (INT): The identifier of the RAID group.

## **raid.cksum.replay.mismatch**

### **Severity**

NOTICE



## Description

This event is issued when we encounter an NVRAM checksum entry which does not match the partner NVRAM shadow checksum entry.

## Corrective Action

(None).

## Syslog Message

Skipping checksum block %ld for disk %s because the checksum 0x%x and shadow checksum 0x%x do not match.

## Parameters

**block\_number** (LONGINT): Checksum block number.

**diskname** (STRING): Name of the disk.

**checksum** (INT): Checksum value.

**shadow\_checksum** (INT): Shadow checksum value.

## raid.cksum.replay.nvgen.mismatch

### Severity

NOTICE

### Description

This event is issued when we encounter an NVRAM generation mismatch for an aggregate during checksum replay. Reparity will start automatically and update checksum blocks involved in recent writes.

### Corrective Action

(None).

### Syslog Message

NVRAM generation mismatch for %s %s%s.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.cksum.replay.plex.missing

### Severity

NOTICE

### Description

This event is issued when we encounter an NVRAM checksum entry for a missing plex during checksum replay. Reparity will start automatically when the plex comes online. At that time, reparity will update checksum blocks involved in recent writes.

### Corrective Action

(None).

## Syslog Message

Cannot find plex with identifier %d in %s %s%s.

## Parameters

**plexid** (INT): The identifier of the plex.  
**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.cksum.replay.plex.unavailable

### Severity

NOTICE

### Description

This event is issued when we encounter an NVRAM checksum entry for an offline or failed plex during checksum replay. Reparity will start automatically when the plex comes online. At that time, reparity will update checksum blocks involved in recent writes.

### Corrective Action

(None).

## Syslog Message

Plex %s in %s %s%s is offline or failed.

## Parameters

**plex** (STRING): The name of the plex.  
**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.cksum.replay.rg.invalid

### Severity

NOTICE

### Description

This event is issued when we encounter an NVRAM checksum entry for a missing RAID group during checksum replay.

### Corrective Action

(None).

## Syslog Message

Invalid RAID group identifier %d in %s %s%s, plex %s.

## Parameters

**rgid** (INT): The identifier of the RAID group.  
**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

**plex** (STRING): The name of the plex.

## **raid.cksum.replay.summary**

### **Severity**

INFORMATIONAL

### **Description**

This event is generated when checksum replay is complete.

### **Corrective Action**

(None).

### **Syslog Message**

Replayed %u checksum blocks.

### **Parameters**

**nblocks** (INT): The number of blocks

## **raid.cksum.replay.vol.clean**

### **Severity**

NOTICE

### **Description**

This event is issued when we encounter an NVRAM checksum entry for a cleanly shutdown aggregate during checksum replay.

### **Corrective Action**

(None).

### **Syslog Message**

%s %s%s was cleanly shut down.

### **Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## **raid.cksum.replay.vol.missing**

### **Severity**

NOTICE

### **Description**

This event is issued when we encounter an NVRAM checksum entry for a missing volume or aggregate during checksum replay. Reparity will start automatically when the aggregate appears again. At that time, reparity will update checksum blocks involved in recent writes.

### **Corrective Action**

(None).

## Syslog Message

Cannot find aggregate with identifier %s.

## Parameters

**volId** (STRING): The identifier of the volume or aggregate.

## raid.cksum.replay.vol.offline

## Severity

NOTICE

## Description

This event is issued when we encounter an NVRAM checksum for an offline aggregate during checksum replay. Reparity will start automatically when the aggregate comes online. At that time, reparity will update checksum blocks involved in recent writes.

## Corrective Action

(None).

## Syslog Message

%s %s%s is not online.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.cksum.unverify.file.block

## Severity

ERROR

## Description

This event is issued when WAFL accesses a block that could not be verified because of medium error on the checksum block.

## Corrective Action

Please restore the corrupted file from snapshot or from the backup media.

## Syslog Message

Reading unverified block on %s %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

## Parameters

**vol\_type** (STRING): Whether the block was being accessed for a flexible volume or an aggregate.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): The name of the volume or aggregate  
**disk\_info** (STRING): Formatted information of the disk  
**blockNum** (LONGINT): The block number  
**ino\_type** (STRING): The inode type  
**fileid** (INT): The file identifier  
**snapid** (INT): The snapshot identifier

**blo** (LONGINT): The block number in the file  
**level** (INT): The level  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.validate.zoned

### Severity

NOTICE

### Description

This event is issued when the aggregate cannot support Zone checksum.

### Corrective Action

(None).

### Syslog Message

Checksums cannot be enabled on %s %s%s due to its disk sizes.

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.cksum.verify.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch in an aggregate when it is being accessed by RAID, and not in response to a WAFL® request; therefore, there is no usable WAFL® context. The expected virtual block Number (VBN)/ disk block number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. Additional messages are logged as the system attempts to perform recovery.

### Corrective Action

(None).

### Syslog Message

Block number mismatch on %s %s%s, %s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu. System will attempt to perform recovery.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Information about the disk object including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Identifier of the disk shelf where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): String representing the name of the model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry in an aggregate when it is being accessed by RAID, and not in response to a WAFL® request; therefore, there is no usable WAFL context. The embedded checksum computed over the checksum entry does not match, indicating the corruption of the checksum entry. Additional messages are logged as the system attempts to perform recovery.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s %s%s, %s, disk block %llu. System will attempt to perform recovery.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Physical disk block number of the disk containing the error.  
**shelf** (STRING): Identifier of the disk shelf where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): String representing the name of the model of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.error

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum error in an aggregate when it is being accessed by RAID, and not in response to a WAFL® request; therefore, there is no usable WAFL® context. Additional messages are logged as the system attempts to perform recovery.

### Corrective Action

(None).

### Syslog Message

%s on %s %s%s, %s, disk block %llu, cpcount %d, buftreeid %d. System will attempt to perform recovery.

### Parameters

**error** (STRING): Name of the error.

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Block number.

**stored\_cpcount** (INT): Stored CP count in the checksum entry.

**stored\_buftreeid** (INT): Stored buftree identifier in the checksum entry.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where disk is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): String representing the name of the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.error.file

### Severity

NOTICE

## Description

This message occurs when the system discovers a checksum error in a block of a file. The computed checksum over the stored data does not match the stored checksum, indicating corruption. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).

## Syslog Message

%s on %s %s%s, %s, disk block %llu, cpcount %d, buftreeid %d, %s inode number %d, snapid %d, file block %llu, level %d.

## Parameters

**error** (STRING): Name of the error.

**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.

**owner** (STRING): Owner of the affected volume or aggregate.

**vol** (STRING): Name of the volume or aggregate.

**disk\_info** (STRING): Formatted information of the disk drive.

**blockNum** (LONGINT): Block number.

**stored\_cpcount** (INT): Stored CP count in the checksum entry.

**stored\_buftreeid** (INT): Stored buftree identifier in the checksum entry.

**ino\_type** (STRING): Inode type.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot(tm) copy identifier.

**bno** (LONGINT): Block number in the file.

**level** (INT): Level.

**shelf** (STRING): Disk shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of the disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.file.bno

## Severity

NOTICE

## Description

This message occurs when the system detects a block number mismatch in a block of a file. The expected virtual block number (VBN)/disk block number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).



## Syslog Message

Block number mismatch on %s %s%s, %s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu, %s inode number %d, snapid %d, file block %llu, level %d.

## Parameters

**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.

**owner** (STRING): Owner of the affected volume or aggregate.

**vol** (STRING): Name of the volume or aggregate.

**disk\_info** (STRING): Formatted information of the disk drive.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_dbn** (LONGINT): Expected physical disk block number.

**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected volume block number.

**ino\_type** (STRING): Inode type.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot(tm) copy identifier.

**bno** (LONGINT): Block number in the file.

**level** (INT): Level.

**shelf** (STRING): Disk shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of the disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.file.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid file block checksum. The embedded checksum does not match the computed checksum, indicating corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions, and other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

### Parameters

**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.

**owner** (STRING): Owner of the affected volume or aggregate.

**vol** (STRING): Name of the volume or aggregate.

**disk\_info** (STRING): Formatted information about the disk drive.

**blockNum** (LONGINT): Physical disk block number that contains the error.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot™ copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Level.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of the disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster™ configuration, indicates the site {local|remote} where the disk is located. For non-MetroCluster configurations, the site is 'local'.  
**physDiskBlockNum** (LONGINT): Physical disk block number of the partitioned disk.

## raid.cksum.verify.file.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry in a block of a file. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

### Parameters

**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot(tm) copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Level.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of the disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.verify.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry for a non-zeroed block in an aggregate when it is being accessed by RAID, and not in response to a WAFL® request; therefore there is no usable WAFL® context. The checksum entry is zeroed, but the corresponding block is not zeroed. Additional messages are logged as the system attempts to perform recovery.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s %s%s, %s, disk block %llu. System will attempt to perform recovery.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical disk block number of the disk containing the error.

**shelf** (STRING): Identifier of the disk shelf where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): String representing the name of the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.wc.blkErr

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum error due to a WAFL® context mismatch. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

## Corrective Action

Please restore the corrupted file from snapshot or from the backup media.

## Syslog Message

Checksum error due to WAFL context mismatch on %s %s%s, %s, disk block %llu, buftree ID %d, %s inode number %d, Snapshot copy ID %d, file block %llu, level %d: checksum context has buftree ID %d, file block %llu, %s.

## Parameters

**vol\_type** (STRING): Whether the file is contained in an aggregate or within a flexible volume.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Block number.  
**buftreeid** (INT): Buftree identifier.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot(tm) copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Indirection level in the WAFL buffer.  
**stored\_buftreeid** (INT): Stored buftree identifier in the checksum entry.  
**stored\_fbn** (LONGINT): Stored file block number in the checksum entry.  
**remaining\_stored\_context** (STRING): String containing CP count, Physical Volume Block Number (PVCN), encrypted flag, encryption key index stored in the checksum entry.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.cksum.wc.sblkErr

### Severity

ALERT

### Description

This message occurs when the system encounters a super block with bad WAFL® context information in its checksum entry.

### Corrective Action

Start wafliroot on the indicated volume or aggregate.

### Syslog Message

Checksum error due to WAFL context mismatch for super block %llu on %s %s%s, %s.

## Parameters

**blockNum** (LONGINT): Block number.  
**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of the disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.zeroed.blkErr

### Severity

NOTICE

### Description

This message occurs when the system discovers that the stored checksum for a block allocated by WAFL® is zero. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum error due to zeroed block on %s %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

## Parameters

**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**blockNum** (LONGINT): Block number.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot(tm) copy identifier.  
**bno** (LONGINT): Block number in the file.  
**level** (INT): Level.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.cksum.zeroed.sblkErr

### Severity

NOTICE

### Description

This message occurs when the system encounters a zeroed super block. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum error due to zeroed super block %llu on %s %s%s, %s.

### Parameters

**blockNum** (LONGINT): Block number.  
**vol\_type** (STRING): Whether the file is contained in an aggregate or a flexible volume.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of the disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**physDiskBlockNum** (LONGINT): Physical disk block number for partitioned disk.

## raid.config events

### raid.config.check.failed

### Severity

ERROR

### Description

This message occurs during a mirror check, when the system discovers that all plexes have failed in an aggregate.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem volume from the system logs. 2. Determine the reason for the plex failures, such as disk drive, disk shelf, or loop failure. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support.

### Syslog Message

%s %s%s: All plexes have failed.

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.config.check.failedPlex

### Severity

ERROR

### Description

This message occurs during a mirror check, when the system discovers that an aggregate has a failed plex.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem volume and failed plex from the system logs. 2. Determine the reason for the plex failure, such as disk drive, disk shelf, or loop failure. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support.

### Syslog Message

Plex %s%s has failed.

### Parameters

**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.config.check.offlinePlex

### Severity

ERROR

### Description

This message occurs during a mirror relationship check when the system discovers that an aggregate has become mirror-degraded due to an offline plex.

### Corrective Action

Bring the plex back online with the "storage aggregate plex online" command.

### Syslog Message

Plex %s%s is offline.

### Parameters

**aggr** (STRING): Name of the aggregate.

**plex** (STRING): Name of the plex.

## **raid.config.disk.bad.label**

### **Severity**

NOTICE

### **Description**

This message occurs when configuration management detects that an unassimilated disk has a label with unexpected or missing data, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

%s has bad or missing data on label.

### **Parameters**

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.config.disk.bad.label.version**

### **Severity**

NOTICE

### **Description**

This message occurs when configuration management detects that an unassimilated disk has a label with an unsupported version, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

%s has an unsupported label version.

### **Parameters**

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number,



vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.failed

### Severity

ERROR

### Description

This message occurs when configuration management has detected that an unassimilated disk has failed, and has permanently removed that disk from service.

### Corrective Action

Wait for the event message raid.disk.unload.done or raid.carrier.remove to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, raid.disk.unload.done will be issued immediately. Otherwise, wait for raid.carrier.remove which indicates the carrier has been fully prepared for removal.

### Syslog Message

%s failed.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.init.failed

### Severity

NOTICE

## Description

This message occurs when configuration management detects that a disk has failed its device initialization sequence, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

## Corrective Action

(None).

## Syslog Message

%s initialization failed.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.labeled.broken

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a filesystem disk was failed under a previous release of Data ONTAP®, and has permanently removed that disk from service. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s was labeled "broken" in a previous release.

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.missing

### Severity

INFORMATIONAL

### Description

This event is issued when configuration management has detected that an unassimilated disk has gone missing from the filer, and has removed that disk from the configuration.

### Corrective Action

(None).

### Syslog Message

%s is missing.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.multipleFail

### Severity

INFORMATIONAL

### Description

This event is generated when we detect additional disk failures during the processing of a failed disk.

### Corrective Action

(None).

### Syslog Message

While processing failed disks, %d additional disks suffered a failure

## Parameters

**ndisks** (INT): The number of additional disks that have failed

## raid.config.disk.not.responding

### Severity

NOTICE

### Description

This message occurs when configuration management detects that an unassimilated disk is temporarily not responding to requests, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s is not responding.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.rawsize.shrank

### Severity

NOTICE

### Description

This message occurs when configuration management detects that an unassimilated disk's stored rawsize is greater than the physical size of the disk, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s shrank.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**failure\_reason** (STRING): Text description of the storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.recovering

### Severity

NOTICE

### Description

This message occurs when configuration management detects an unassimilated disk that has failed to respond to requests. The disk is temporarily removed from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Attempting to bring %s back into service.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**failure\_reason** (STRING): Text description of the storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.disk.tradvol

### Severity

NOTICE

## Description

This message occurs when configuration management detects that the disk belongs to a traditional volume which is not supported on clustered Data ONTAP® systems. The disk is removed from service.

## Corrective Action

Install a version of Data ONTAP® that supports the traditional volume, or remove the traditional volume disks.

## Syslog Message

The label on %s indicates that the disk belongs to a traditional volume which is not supported on clustered Data ONTAP® systems.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.admin.failed

### Severity

INFORMATIONAL

### Description

This event is issued after the "disk fail -i" command is used to administratively fail a filesystem disk, and configuration management has permanently removed that disk from service.

### Corrective Action

(None).

### Syslog Message

File system %s is being failed by administrator.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.config.filesystem.disk.admin.failed.after.copy**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the "disk fail" command is used to administratively fail a filesystem disk. The disk was copied to a replacement and configuration management has permanently removed that disk from service.

### **Corrective Action**

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### **Syslog Message**

File system %s is being failed after it was successfully copied to a replacement.

### **Parameters**

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.config.filesystem.disk.bad.label**

### **Severity**

NOTICE

### **Description**

This message occurs when configuration management detects that a filesystem disk has a label with unexpected or missing data, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

## Syslog Message

File system %s has bad or missing data on label.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.bad.label.version

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a filesystem disk has a label with an unsupported version, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

File system %s has an unsupported label version.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.



## raid.config.filesystem.disk.failed

### Severity

ERROR

### Description

This message occurs when configuration management has detected that a filesystem disk has failed, and has permanently removed that disk from service.

### Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### Syslog Message

File system %s failed.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.failed.after.copy

### Severity

ERROR

### Description

This message occurs when a disk reported a predictive failure. Data ONTAP used Rapid RAID Recovery to copy the disk to a replacement. Copying was successful and now the disk is permanently removed from service.

### Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### Syslog Message

File system %s is being failed after it was successfully copied to a replacement.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**failure\_reason** (STRING): A text description of storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.missing

### Severity

INFORMATIONAL

### Description

This event is issued when configuration management has detected that a filesystem disk has gone missing from the filer, and has removed that disk from the configuration.

### Corrective Action

(None).

### Syslog Message

File system %s is missing.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**failure\_reason** (STRING): A text description of storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.not.responding

### Severity

NOTICE

## Description

This message occurs when configuration management detects that a filesystem disk is temporarily not responding to requests, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

## Corrective Action

(None).

## Syslog Message

File system %s is not responding.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.rawsize.shrank

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a filesystem disk's stored rawsize is greater than the physical size of the disk, and temporarily removes that disk from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

File system %s shrank.

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.disk.recovering

### Severity

NOTICE

### Description

This message occurs when configuration management detects a filesystem disk that has failed to respond to requests. The disk is temporarily removed from service. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Attempting to bring file system %s back into service.

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.filesystem.lun.resized

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a filesystem gateway LUN's stored rawsize does not match the physical size of the LUN. Resizing of a filesystem gateway LUN is not supported. Only gateway LUNs that are spare can be resized. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

File system %s resized.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.online.notComp

### Severity

ERROR

### Description

This message occurs when a volume or aggregate cannot be brought online because SnapLock® is not supported on this version of Data ONTAP®.

### Corrective Action

Install a version of Data ONTAP that is supported with SnapLock, available from your vendor.

### Syslog Message

%s %s%s is being taken offline. SnapLock is not supported on this version of Data ONTAP.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.config.online.notCompressComp

### Severity

ERROR

### Description

This message occurs when a compressed virtual volume cannot be brought online because compressed volumes have adaptive storage efficiency which is not supported in this release of Data ONTAP®.

### Corrective Action

Install a version of Data ONTAP that supports adaptive storage efficiency.

## Syslog Message

Volume '%s%s' is being taken offline because it is compressed and the container uses adaptive storage efficiency which is not supported in this version of Data ONTAP.

## Parameters

**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.

## raid.config.online.req.bad.fsid

### Severity

NOTICE

### Description

This message occurs when an aggregate cannot be mounted because it has an invalid file system identifier (FSID) or the same FSID as another aggregate.

### Corrective Action

(None).

## Syslog Message

%s %s%s has a duplicate or invalid FSID.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.config.online.req.badvers

### Severity

ERROR

### Description

This message occurs when an aggregate that appeared to be mountable could not be mounted because the WAFL® version is invalid.

### Corrective Action

Install a new version of Data ONTAP® because the aggregate is from a later version of software.

## Syslog Message

This version of Data ONTAP does not recognize the filesystem of %s %s%s. The %s is probably from a later version of the software and is being left unmounted.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate object that could not be mounted.  
**vol\_type2** (STRING): Repeat volume type for the syslog message.

## raid.config.online.req.corrupt4

### Severity

ERROR

### Description

This event occurs when dynamic assimilation produced a writable aggregate which could not be mounted because it has been marked as having an inconsistent file system.

### Corrective Action

Run wafliiron for the aggregate.

### Syslog Message

%s %s%s is inconsistent. Mounting it prior to running wafliiron raises the risk of further file system inconsistencies and data corruption.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that could not be mounted.

## raid.config.online.req.failed

### Severity

NOTICE

### Description

This message occurs when an attempt to mount an aggregate fails, even though it appeared to be mountable.

### Corrective Action

(None).

### Syslog Message

Unable to mount %s %s%s.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that could not be mounted.

## raid.config.online.req.inconsist

### Severity

ERROR

### Description

This event occurs when dynamic assimilation produced an aggregate which appeared to be mountable, but it could not be mounted because the WAFL magic number in the volume information block was incorrect. This is caused by an incomplete aggregate copy or snapmirror initialization.

### Corrective Action

Repeat aggr copy or snapmirror initialize that was not completed.

### Syslog Message

%s %s%s was left in an inconsistent state by an aborted vol copy or aggr copy copy or an aborted snapmirror initial (level 0) transfer. In order to bring it online, you must either destroy and re-create it, or complete an initial snapmirror transfer or vol copy or aggr copy.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that could not be mounted.

## raid.config.online.req.invalid

### Severity

NOTICE

### Description

This message occurs when an aggregate that appeared to be mountable could not be mounted because neither volume info block was valid.

### Corrective Action

(None).

### Syslog Message

Neither volume information block of %s %s%s is valid.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that could not be mounted.

## raid.config.online.req.mt.iv

### Severity

ERROR

### Description

This message occurs when WAFL® attempts to take over a mirror volume that does not have a valid mirror type. This might be caused by upgrading to the current version of Data ONTAP®, which requires the mirror-type attribute, from a previous version that does not.

### Corrective Action

You cannot fix this condition by file system repair mechanisms such as wafiron. The data is not lost, but it is inaccessible as a mirror volume with invalid mirror information, and it cannot be brought online or taken offline. To reclaim the space you can destroy the mirror volume, but this results in the loss of all data on the mirror volume. If you choose to destroy the mirror volume anyway, the command to do so is as follows: 'volume delete -vserver vserver\_name -volume vol\_name'.



## Syslog Message

The mirror type associated with the mirror volume %s %s%s is invalid.

## Parameters

**vol\_type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the affected volume.

**vol** (STRING): Name of the volume.

## raid.config.online.req.nospace

### Severity

ERROR

### Description

This event occurs when there is insufficient free space in the containing aggregate to online a flexible volume and apply the flexible volumes space guarantees.

### Corrective Action

This error can be corrected by making more space available in the aggregate.

## Syslog Message

Containing aggregate does not have enough space to allow %s to be onlined.

## Parameters

**vol** (STRING): The name of the flexible volume which could not be onlined.

## raid.config.online.req.ok

### Severity

NOTICE

### Description

This message occurs when dynamic assimilation produced an aggregate which was successfully mounted.

### Corrective Action

(None).

## Syslog Message

%s %s%s is mounted. Time taken to mount the aggregate was %llu milliseconds.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that was mounted.

**time** (LONGINT): Time taken to mount the aggregate.

## raid.config.online.req.unsup

### Severity

EMERGENCY

## Description

This message occurs when the system determines that the deployment in which this aggregate was created does not match the current (clustered or nonclustered) deployment of Data ONTAP®.

## Corrective Action

Bring the aggregate online on a node that matches the clustered or nonclustered deployment in which it was created.

## Syslog Message

This deployment of Data ONTAP does not support the file system of '%s %s'. It is being taken offline.

## Parameters

**type** (STRING): Volume type.

**vol** (STRING): Name of the aggregate.

## raid.config.raidsize.change

### Severity

NOTICE

### Description

This message occurs when a change to the raidsize parameter is made to match the change made to the raidtype parameter for a given volume or aggregate.

### Corrective Action

If the new raidsize is not appropriate, change it with the "storage aggregate modify -aggregate aggregate\_name -maxraidsize new\_size" command. For cache raidsize change, change it with the "storage aggregate modify -aggregate aggregate\_name -cache\_raid\_group\_size new\_size" command.

### Syslog Message

%s %s%s:%sraidsize is adjusted from %d to %d after changing raidtype.

### Parameters

**vol\_type** (STRING): Volume type.

**host** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**cache\_string** (STRING): String to indicate if its cache raidsize change. If this is a raidsize change for cache (SSD) RG then the string "cache" will be printed, it is empty otherwise.

**old\_size** (INT): Previous raidsize value.

**new\_size** (INT): New raidsize value.

## raid.config.rg.size.reminder

### Severity

ERROR

### Description

Size limits for raid4 groups are significantly reduced in Data ONTAP® 6.5. Very large raid4 groups do not provide adequate protection from data loss and they are not supported any more. Volumes with large raidgroups should be converted to raid\_dp. This message occurs daily as part of system check if some raid groups are too large for raid4 raidtype. That can happen after upgrade from an older version of Data ONTAP with more lax raidsize restrictions. It can also be the result of changing raidtype from raid\_dp to

raid4. That should be used only to revert to an older version of Data ONTAP. In all other cases, change raidtype of volumes with such large raid groups to raid\_dp.

### Corrective Action

Change raidtype of the aggregate to raid\_dp: "storage aggregate modify -aggregate aggregate\_name -raidtype raid\_dp"

### Syslog Message

Volume %s%s has a raid group larger than the limit for raid4 raidtype (max. raid group size: %d, limit is %d). Convert the volume to raid\_dp.

### Parameters

**host** (STRING): Owner of the volume.  
**vol** (STRING): Volume name.  
**rg\_size** (INT): Size of the largest raid group in the volume.  
**limit** (INT): Size limit for raid4 group.

## raid.config.spare.disk.admin.removed

### Severity

INFORMATIONAL

### Description

This event is issued when the "disk remove" command is used to administratively fail a spare disk, and configuration management has temporarily removed that disk from service.

### Corrective Action

(None).

### Syslog Message

Spare %s is being removed by administrator.

### Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**failure\_reason** (STRING): A text description of storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.bad.label

### Severity

NOTICE

## Description

This message occurs when configuration management detects that a spare disk has a label with unexpected or missing data, and has temporarily removed that disk from service.

## Corrective Action

(None).

## Syslog Message

Spare %s has bad label.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.bad.label.version

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a spare disk has a label with an unsupported version, and has temporarily removed that disk from service.

### Corrective Action

(None).

### Syslog Message

Spare %s has an unsupported label version.

### Parameters

**disk\_info** (STRING): Formatted information of the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.failed

### Severity

ERROR

### Description

This message occurs when configuration management has detected that a spare disk has failed, and has permanently removed that disk from service.

### Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### Syslog Message

Spare %s failed.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.missing

### Severity

INFORMATIONAL

### Description

This event is issued when configuration management has detected that a spare disk has gone missing from the filer, and has removed that disk from the configuration.

### Corrective Action

(None).

### Syslog Message

Spare %s is missing.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**failure\_reason** (STRING): A text description of storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.not.responding

### Severity

NOTICE

### Description

This message occurs when configuration management detects that a spare disk is not responding to requests, and has temporarily removed that disk from service.

### Corrective Action

(None).

### Syslog Message

Spare %s is not responding.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**failure\_reason** (STRING): Text description of the storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.rawsize.shrank

### Severity

NOTICE

## Description

This message occurs when configuration management detects that a spare disk's stored rawsize is greater than the physical size of the disk, and has temporarily removed that disk from service.

## Corrective Action

(None).

## Syslog Message

Spare %s shrank.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.config.spare.disk.recovering

### Severity

NOTICE

### Description

This message occurs when configuration management detects a spare disk that has failed to respond to requests. The disk is temporarily removed from service. Attempts will be made to automatically bring this disk back into service. Unless the disk continues to exhibit the same problem, recovery should be complete within one hour. Otherwise, a follow-on notification will be sent within five hours to signal that the disk is being permanently removed from service.

### Corrective Action

(None).

### Syslog Message

Attempting to bring spare %s back into service.

### Parameters

**disk\_info** (STRING): Formatted information of the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.config.vol.destroyed**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when an aggregate is destroyed.

### **Corrective Action**

(None).

### **Syslog Message**

%s '%s%s' destroyed.

### **Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**volumeName** (STRING): Name of the aggregate object.

## **raid.config.vol.renamed**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when an aggregate is renamed.

### **Corrective Action**

(None).

### **Syslog Message**

%s '%s%s' renamed to '%s%s'.

### **Parameters**

**vol\_type** (STRING): Volume type.

**oldOwner** (STRING): Owner of the aggregate.

**oldVolumeName** (STRING): The original name of the object.

**newOwner** (STRING): String indicating the owner of the affected object. This is exactly the same as the oldOwner parameter (or should be).

**newVolumeName** (STRING): The new name of the object.

## **raid.config.vol.set.raid.cv**

### **Severity**

NOTICE



## Description

This event is issued when an aggregate with block checksums is changing its checksum verification flag value.

## Corrective Action

(None).

## Syslog Message

%s %s%s: checksum verification %s.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**volumeName** (STRING): Name of the aggregate object.  
**status** (STRING): The status of the command.

# raid.data events

## raid.data.lw.blkErr

### Severity

NOTICE

### Description

This message occurs when a WAFL® generated read detects a stale block due to a lost write. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Bad data detected on %s%s, disk block (DBN) %llu, aggregate block (VBN) %llu.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number that contains the error.  
**blockNum** (LONGINT): Physical block number containing the error.  
**vbn** (LONGINT): Volume block number containing the error.

## raid.data.lw.recFail

### Severity

ALERT

### Description

This message occurs when raid recommends a drive for failure because it detected more than the allowed number of lost-write errors on the disk. These errors are caused by a malfunction in the disk.

## Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

## Syslog Message

Disk recommended for failure for exceeding threshold of %d lost-write error%s on %s%s.

## Parameters

**threshold** (INT): The threshold for checksum errors due to wafI context mismatches  
**suffix** (STRING): The suffix used to report in syslog  
**owner** (STRING): String indicating the owner of the affected volume.  
**disk\_info** (STRING): The name of the disk.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.data.ws.blkErr

### Severity

NOTICE

### Description

This message occurs when RAID detects a stale data block using WAFL® sanity check. The bad block is then reconstructed from parity and the remaining data disks. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

WAFL sanity check detected bad data on volume %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d.

### Parameters

**owner** (STRING): Owner of the affected volume.  
**vol** (STRING): Name of the volume.  
**disk\_info** (STRING): Formatted information about the disk.  
**blockNum** (LONGINT): Block number.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): File identifier.  
**snapid** (INT): Snapshot™ identifier.  
**fileBlockNum** (LONGINT): Block number in the file.

**level** (INT): Indirection level.

## **raid.data.ws.blkErr.ignore**

### **Severity**

NOTICE

### **Description**

This message occurs when WAFL® sanity check detects a stale data block while wafiron is running. RAID ignores the error because wafiron fixes it. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL sanity check detected bad data on volume %s%s, %s, disk block %llu, %s inode number %d, snapid %d, file block %llu, level %d. RAID is ignoring this error because wafiron is already running.

### **Parameters**

**owner** (STRING): String indicating the owner of the affected volume.

**vol** (STRING): Name of the volume.

**disk\_info** (STRING): Formatted information about the disk.

**blockNum** (LONGINT): Block number.

**ino\_type** (STRING): Inode type.

**fileid** (INT): File identifier.

**snapid** (INT): Snapshot™ identifier.

**fileBlockNum** (LONGINT): Block number in the file.

**level** (INT): Indirection level.

## **raid.data.ws.sblkErr**

### **Severity**

NOTICE

### **Description**

This message occurs when a WAFL® sanity check detects a stale superblock. The bad superblock is reconstructed from parity and the remaining data disks. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL sanity check detected bad superblock on volume %s%s, %s, disk block (DBN) %llu, aggregate block (VBN) %llu.

### **Parameters**

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the volume.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.

**blockNum** (LONGINT): Block number.  
**vbn** (LONGINT): Volume block number.

## raid.data.ws.sblkErr.ignore

### Severity

NOTICE

### Description

This message occurs when WAFL® sanity check detects a stale superblock while wafiron is running. RAID ignores the error because wafiron tries to fix it. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

WAFL sanity check detected bad superblock on volume %s%s, %s, block %llu. RAID is ignoring this error because wafiron is already running.

### Parameters

**owner** (STRING): String indicating the owner of the affected volume.  
**vol** (STRING): Name of the volume.  
**disk\_info** (STRING): Formatted information of the disk.  
**blockNum** (LONGINT): Block number.

## raid.debug events

### raid.debug

### Severity

INFORMATIONAL

### Description

The SCSI debug event.

### Corrective Action

(None).

### Syslog Message

%s

### Parameters

**debug\_string** (STRING): The debug statement that occurred.

## raid.disable events

### raid.disable.remote.nvlog

**Severity**

INFORMATIONAL

**Description**

This event is issued when RAID sub system disables recording in the partner's NVRAM entries for the log based file system.

**Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

**Syslog Message**

Disabled Remote NVRAM logging.

**Parameters**

(None).

## raid.disk.events

### raid.disk.capacity.info

**Severity**

NOTICE

**Description**

This event is issued to inform what the sum of all disk sizes is as well as what system capacity limit is.

**Corrective Action**

(Call support).

**Syslog Message**

System capacity visible is %s of the %s maximum.

**Parameters**

**actual** (STRING): Actual capacity of system

**limit** (STRING): Capacity limit of system

### raid.disk.cksum.type.changed

**Severity**

INFORMATIONAL

**Description**

This message occurs when a disk's checksum type changes from either block to zoned, or zoned to block. It might occur, for example, due to a disk firmware update or a user request.

**Corrective Action**

(None).

**Syslog Message**

Disk %s checksum type was changed from %s to %s.

## Parameters

**disk\_info** (STRING): Formatted disk information  
**old\_type** (STRING): Old checksum type. "block" or "zoned".  
**new\_type** (STRING): New checksum type. "block" or "zoned".  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.failed.unknownPool

### Severity

INFORMATIONAL

### Description

This event is issued when a disk is failed but is not in an aggregate or a known disk pool.

### Corrective Action

(None).

### Syslog Message

Failed %s without a tree but with a unknown pool as a parent.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.fast.zero.done

### Severity

NOTICE

### Description

This message occurs when a fast zeroing operation is completed on a disk.

## Corrective Action

(None).

## Syslog Message

%s : disk zeroing complete (0x%llx).

## Parameters

**disk\_info** (STRING): Formatted information about the disk which has completed its zeroing operation.

**fz\_signature** (LONGINTEX): Random generated unique fast zero signature.

**fz\_signature\_old** (LONGINTEX): Random generated unique fast zero signature, previous one.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of disk in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.home.id.changed

### Severity

INFORMATIONAL

### Description

This event is issued when a change in a disk's home id is detected, independent of any disk change notification event.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk\_info** (STRING): Formatted disk information

**old\_home\_id** (STRING): Previous home host.

**new\_home\_id** (STRING): New home host.

**old\_home\_sysid** (LONGINT): Previous home sysid.

**new\_home\_sysid** (LONGINT): New home sysid.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.illegalAttach

### Severity

ALERT

### Description

This message occurs when ONTAP detects that a disk has been incorrectly attached to the system.

### Corrective Action

Check disk shelf cabling and verify it meets SyncMirror requirements and wiring rules for this system.

### Syslog Message

%s is illegally attached. %s

### Parameters

**disk\_info** (STRING): Name of the disk.

**err\_string** (STRING): Illegal Attach error.

**instancelid** (INT): Instance identifier.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of disk in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.inserted

### Severity

INFORMATIONAL

### Description

This event is issued when a disk is inserted into the system.

### Corrective Action

(None).

### Syslog Message

%s has been inserted into the system

### Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk



**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.io.bypass

### Severity

INFORMATIONAL

### Description

This event is issued when a read error is detected but ignored because the aggregate is in media error bypass mode.

### Corrective Action

(None).

### Syslog Message

media error bypass enabled, ignoring read error on disk block %llu of %s

### Parameters

**blockNum** (LONGINT): The disk block number  
**disk\_info** (STRING): The formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.io.toFailedDisk

### Severity

ERROR

### Description

This event is issued when an attempt is made to I/O a failed disk.

### Corrective Action

In maintenance mode, remove ownership of this disk and assign the correct ownership

## Syslog Message

DBG: %s is attempting I/O to failed %s

## Parameters

**who** (STRING): Module attempting I/O  
**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.abort

### Severity

NOTICE

### Description

This event occurs when the administrator aborts maintenance testing of a disk or when an error that is unrelated to the disk occurs. The disk is failed because testing could not be completed.

### Corrective Action

(None).

## Syslog Message

%s tests were aborted.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**start\_code** (STRING): Start code for internal support use  
**end\_code** (STRING): End code for internal support use  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk  
**disk\_rpm** (STRING): Rotational speed of disk, in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.admin.abort

## Severity

NOTICE

## Description

This event occurs when administrator-initiated maintenance testing of a disk is aborted by the administrator or when an error that is unrelated to the disk occurs. The disk is made available as a spare.

## Corrective Action

(None).

## Syslog Message

%s tests were aborted.

## Parameters

**disk\_info** (STRING): Formatted information of the disk

**end\_code** (STRING): End code for internal support use

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where the disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk

**disk\_rpm** (STRING): Rotational speed of disk, in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.admin.done

### Severity

NOTICE

### Description

This event occurs when administrator- initiated maintenance testing of a disk is completed successfully. The disk is made available as a spare.

### Corrective Action

(None).

### Syslog Message

%s tests were completed successfully.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where the disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk

**disk\_rpm** (STRING): Rotational speed of disk, in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.admin.failed

### Severity

NOTICE

### Description

This message occurs when a disk fails one or more maintenance tests. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s tests failed.

### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**end\_code** (STRING): End code for internal support use.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.admin.start

### Severity

NOTICE

### Description

This event occurs when the administrator starts maintenance testing on a disk.

### Corrective Action

(None).

### Syslog Message

%s will be tested.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk  
**disk\_rpm** (STRING): Rotational speed of disk, in RPM  
**start\_reason** (STRING): This states that the administrator started maintenance testing on a disk.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.bypassed

### Severity

INFORMATIONAL

### Description

This message occurs when a disk that has an error bypasses maintenance testing, even though the error typically would cause the disk to enter the maintenance center. The reason might be that the maintenance center is disabled, that there are not enough spare disks available in the system, or that there are already the maximum number of disks in the maintenance center.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk\_info** (STRING): Formatted information of the disk.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of the disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**bypass\_reason** (STRING): Reason for bypassing maintenance testing on the disk.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.done

### Severity

NOTICE

### Description

This event occurs when maintenance testing of a disk is completed successfully. The disk is made available as a spare.

## Corrective Action

(None).

## Syslog Message

%s tests were completed successfully.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**start\_code** (STRING): Start code for internal support use  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where the disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk  
**disk\_rpm** (STRING): Rotational speed of disk, in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.exceed

### Severity

NOTICE

### Description

This message occurs when a disk exceeds the number of times it is allowed to be tested. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s exceeded the number of allowed entries to maintenance testing.

### Parameters

**disk\_info** (STRING): Formatted information about the disk.  
**start\_code** (STRING): Start code for internal support use.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.failed

### Severity

NOTICE

### Description

This message occurs when a disk fails one or more maintenance tests. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s tests failed.

### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**start\_code** (STRING): Start code for internal support use.

**end\_code** (STRING): End code for internal support use.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.maint.start

### Severity

NOTICE

### Description

This event occurs when a disk begins maintenance testing.

### Corrective Action

(None).

### Syslog Message

%s will be tested.

### Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where the disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk  
**disk\_rpm** (STRING): Rotational speed of disk, in RPM  
**start\_reason** (STRING): The reason for starting maintenance testing on the disk.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.mcc.mismatch

### Severity

ERROR

### Description

This message occurs when a disk with a non-zero disaster recovery (DR) home ID is illegally attached to a system that is not configured for MetroCluster(tm) operation.

### Corrective Action

To reuse this disk on a non MetroCluster enabled system, use the "storage disk removeowner" command, followed by the "storage disk assign" command.

### Syslog Message

The disk %s was previously owned by a MetroCluster enabled system. This system is not configured for MetroCluster operations.

### Parameters

**disk\_info** (STRING): Name of the disk.  
**instanceid** (INT): Instance identifier.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the disk's DR home owner.

## raid.disk.missing

### Severity

INFORMATIONAL

### Description

This event is issued either when a disk is pulled from the system or when all paths to it have disappeared.

### Corrective Action

(None).



## Syslog Message

%s is missing from the system

## Parameters

**disk\_info** (STRING): Formatted disk information

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.offline

### Severity

NOTICE

### Description

This event is issued when offlining a drive

### Corrective Action

(None).

### Syslog Message

Marking %s%s offline.

### Parameters

**owner** (STRING): The owner of the drive

**disk\_info** (STRING): Formatted information of the disk

## raid.disk.online

### Severity

NOTICE

### Description

This event is issued when onlining an offline drive.

### Corrective Action

(None).

### Syslog Message

Onlining %s%s.

### Parameters

**owner** (STRING): The owner of the drive

**disk\_info** (STRING): Formatted information of the disk

## **raid.disk.online.fail**

### **Severity**

NOTICE

### **Description**

This event is issued if the drive failed while it was offline. The drive will be moved to broken pool as part of the online process.

### **Corrective Action**

(None).

### **Syslog Message**

Failing %s%s during online.

### **Parameters**

**owner** (STRING): The owner of the drive

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.disk.owner.change.fail**

### **Severity**

EMERGENCY

### **Description**

This message occurs when an ownership change request initiated by RAID, for either spare disks or the disks of an aggregate, fails due to an error from SANOWN.

### **Corrective Action**

The system automatically retries changing ownership of the disks. If the condition persists, contact NetApp technical support.

### **Syslog Message**

RAID initiated ownership change request failed (aggregate: %s error: %s). This request will be retried.

### **Parameters**

**aggregate** (STRING): Empty string for spare disks, or name of the aggregate for filesystem disks.

**reason** (STRING): Reason for failure.

**retry\_count** (INT): Number of retry attempts.

**disk\_count** (INT): Number of disks in the ownership change request.

## **raid.disk.owner.changed**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued when change in a disk's owner is detected, independent of any disk change notification event.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**disk\_info** (STRING): Formatted disk information

**old\_owner** (STRING): Previous disk owner.

**new\_owner** (STRING): New disk owner.

**old\_owner\_id** (LONGINT): Old disk owner ID.

**new\_owner\_id** (LONGINT): New disk owner ID.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.disk.ownership.change**

### **Severity**

NOTICE

### **Description**

This message occurs when a non-filesystem disk is prepared for ownership change by administrator forcibly.

### **Corrective Action**

(None).

### **Syslog Message**

Ownership of %s, is being changed by administrator forcibly.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.pool.changed

### Severity

INFORMATIONAL

### Description

This message occurs when a change in a disk's current pool is detected, independent of any disk change notification event.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**disk\_info** (STRING): Formatted disk information.  
**old\_pool** (STRING): Previous pool for this disk.  
**new\_pool** (STRING): New pool for this disk.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.predictiveFailure

### Severity

ERROR

## Description

This message occurs when some condition of the disk indicates that it may fail soon. It is first marked as prefailed. If an appropriate spare is available, it will be selected for Rapid RAID Recovery. In that process, the prefailed disk will be copied to the spare. The filer will stop using the prefailed disk when copying is done. At that time, the disk may be removed and replaced.

## Corrective Action

The raid.rg.diskcopy.done event is issued when the disk copy is complete. After that wait for the event message raid.disk.unload.done or raid.carrier.remove to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, raid.disk.unload.done will be issued immediately. Otherwise, wait for raid.carrier.remove which indicates the carrier has been fully prepared for removal.

## Syslog Message

%s reported a predictive failure and it is prefailed; it will be copied to a spare and failed

## Parameters

**disk\_info** (STRING): Formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.read.abort.detect

### Severity

NOTICE

### Description

This message occurs when a disk encounters an abort error during a read operation. The system retries the disk read request and gradually takes the disk out of service by copying the content to a replacement disk.

### Corrective Action

(None).

### Syslog Message

Read abort error encountered on %s, disk rebuild assist mode %s, block #%llu, length = %u; I/O will be retried.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**disk\_ra\_mode** (STRING): Disk rebuild assist mode indicating whether mode is enabled.

**blockNum** (LONGINT): Physical disk block number of the first block in this I/O chain.

**length** (INT): Number of disk blocks in this I/O chain.

**error\_code** (INT): Error code returned by the disk driver.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Disk firmware revision number.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.read.abort.recons

### Severity

NOTICE

### Description

This message occurs when a disk encounters an abort error during a read operation, and a subsequent attempt to retry the read request fails. The system services the read request by reconstructing the data from parity data. The system gradually takes the disk out of service by copying the content to a replacement disk.

### Corrective Action

(None).

### Syslog Message

Read retry failed on %s, block #%llu.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.  
**blockNum** (LONGINT): Physical disk block number on which the read operation failed.  
**error\_code** (INT): Error code returned by the disk driver.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Disk firmware revision number.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.replace.job.done

### Severity

NOTICE

**Description**

This message occurs when a disk replace job finishes successfully.

**Corrective Action**

(None).

**Syslog Message**

Disk %s replaced with disk %s.

**Parameters**

**srcName** (STRING): Name of the source disk.

**dstName** (STRING): Name of the destination disk.

**raid.disk.replace.job.failed****Severity**

ERROR

**Description**

This message occurs when a disk replace job fails.

**Corrective Action**

Take any corrective action indicated by the specific reason for the failure and then retry the disk replace operation. If the condition persists, contact NetApp technical support.

**Syslog Message**

Failed to replace disk %s with disk %s. Reason: %s.

**Parameters**

**srcName** (STRING): Name of the source disk.

**dstName** (STRING): Name of the destination disk.

**reason** (STRING): Reason for failure.

**raid.disk.replace.job.start****Severity**

NOTICE

**Description**

This message occurs when a job to replace a shared disk with a spare disk begins. The operation can be triggered either by an administrator or automatically by Data ONTAP®. A couple of such scenarios follows:  
a) There is a mix of drives from different SyncMirror® pools in the same plex of a mirrored aggregate. b) An aggregate is using a partitioned disk on an external shelf when there is a spare disk available on an internal shelf.

**Corrective Action**

(None).

**Syslog Message**

Starting disk replacement of disk %s with disk %s.

## Parameters

**srcName** (STRING): Name of the source disk.  
**dstName** (STRING): Name of the destination disk.

## raid.disk.replace.job.stop

### Severity

NOTICE

### Description

This message occurs when a disk replace job is stopped by an administrator.

### Corrective Action

(None).

### Syslog Message

Stopping disk replacement of disk %s with disk %s.

## Parameters

**srcName** (STRING): Name of the source disk.  
**dstName** (STRING): Name of the destination disk.

## raid.disk.sanitization.aborted

### Severity

INFORMATIONAL

### Description

This event is issued when disk sanitization is aborted using the "disk sanitize" abort" command.

### Corrective Action

(None).

### Syslog Message

Sanitization aborted for %s

## Parameters

**disk\_info** (STRING): Formatted information of the disk object  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'



## raid.disk.sanitization.failed

### Severity

INFORMATIONAL

### Description

This event is issued if disk sanitization fails due to a fatal error.

### Corrective Action

(None).

### Syslog Message

Sanitization failed for %s

### Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, the site is 'Local'.

## raid.disk.timeout.force.recovery

### Severity

NOTICE

### Description

This event is issued when a disk continues to exhibit symptoms of long response times (timeouts) while it is in the process of being replaced via a copy. In this case, a forced recovery is initiated by removed the disk from the raid group. A complete rewrite of the disk is expected to help recover bad media patches.

### Corrective Action

(None).

### Syslog Message

Forcing recovery of %s%s as a result of exceeding timeout recovery threshold of %d.

### Parameters

**owner** (STRING): String indicating the owner of the affected volume.

**disk\_info** (STRING): The name of the disk.

**threshold** (INT): Threshold for the maximum number of timeouts which can be tolerated while a disk is in the process of being replaced.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.timeout.recommend.recovery

### Severity

NOTICE

### Description

This event is issued when a disk repeatedly exhibits symptoms of long response times (timeouts) due to bad media patches. The disk is recommended for a full recovery as a result of crossing timeout recovery thresholds. The contents of the disk will be copied to a replacement. A complete rewrite of the disk is expected to help recover bad media patches.

### Corrective Action

(None).

### Syslog Message

%s%s recommended for recovery as a result of exceeding timeout recovery threshold of %d.

### Parameters

**owner** (STRING): String indicating the owner of the affected volume.  
**disk\_info** (STRING): The name of the disk.  
**threshold** (INT): Threshold for the maximum number of timeout recovery operations above which the disk is recommended for recovery.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.timeout.recovery.invalid.range

### Severity

NOTICE

### Description

This event is issued if an invalid blockNum range was reported during disk timeout recovery

## Corrective Action

(None).

## Syslog Message

Invalid block range reported for %s during disk timeout recovery : start %llu, end %llu

## Parameters

**disk\_info** (STRING): Formatted information of the disk for which the invalid blockNum range was reported.

**start** (LONGINT): Starting disk block number for the recovery range

**end** (LONGINT): Ending disk block number for the recovery range

## raid.disk.timeout.recovery.read.err

### Severity

NOTICE

### Description

This event is issued when timeout recovery io detects a read error

### Corrective Action

(None).

## Syslog Message

Read error on %s%s, block #%llu during aggressive timeout recovery

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.tooBig.all.reminder

### Severity

ERROR

### Description

This message occurs when the sum of disk sizes of all disks connected in the system exceed the maximum capacity supported for this platform.

## Corrective Action

Contact NetApp technical support. For information about correcting the problem, search for the "raid.disk.tooBig.all.reminder" keyword on the Knowledgebase of the NetApp Support Site.

## Syslog Message

System capacity %s exceeds the %s maximum supported disk capacity.

## Parameters

**actual** (STRING): Actual total capacity of disks currently connected to the system.

**limit** (STRING): Maximum supported disk capacity limit of the system.

## raid.disk.unfail.done

### Severity

INFORMATIONAL

### Description

This event is issued when a failed disk has been unfailed using the "disk unfail" command, and returned to service as a spare.

### Corrective Action

(None).

### Syslog Message

%s unfailed, and is now a spare

### Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.unfail.error

### Severity

INFORMATIONAL

### Description

This event is issued when an error is encountered while attempting to unfailed a disk.

### Corrective Action

(None).

## Syslog Message

%s could not be unfailed

## Parameters

**disk\_info** (STRING): Formatted information of the disk object

**error\_str** (STRING): A raid error string, describing the error encountered during disk unfailed.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.unfailed.reassim

### Severity

INFORMATIONAL

### Description

This event occurs when a failed disk was unfailed using the "disk unfailed" command, and is in process of being reassimilated.

### Corrective Action

(None).

## Syslog Message

%s was unfailed, and is now being reassimilated

## Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.unload.done

### Severity

INFORMATIONAL

## Description

This event is issued when a failed or removed disk has been spun down by the driver.

## Corrective Action

It is safe to pull the disk from the shelf.

## Syslog Message

Unload of %s has completed successfully

## Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.verify.abort.detec

### Severity

NOTICE

### Description

This message occurs when a disk encounters an abort error during a background media verify operation. The system gradually takes the disk out of service by copying the content to a replacement disk.

### Corrective Action

(None).

### Syslog Message

Media verify abort error encountered on %s, block #%llu, length = %llu.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**blockNum** (LONGINT): Physical disk block number of the first block in this I/O chain.

**length** (LONGINT): Number of disk blocks in this I/O chain.

**error\_code** (INT): Error code returned by the disk driver.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.zero.aborted

### Severity

NOTICE

### Description

This message occurs when disk zeroing has been aborted. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s : disk zeroing aborted (%s).

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.

**reason** (STRING): Reason for abort.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disk.zero.done

### Severity

NOTICE

### Description

This message occurs when a disk zeroing operation completes.

### Corrective Action

(None).

### Syslog Message

%s : disk zeroing complete

### Parameters

**disk\_info** (STRING): Formatted information of disk object that has completed its zeroing operation

**duration** (INT): The amount of time the zeroing operation required, in seconds

**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.disk.zero.done.failed.die**

### **Severity**

NOTICE

### **Description**

This message occurs when a disk zeroing operation completes.

### **Corrective Action**

(None).

### **Syslog Message**

%s : disk zeroing for failed die complete

### **Parameters**

**disk\_info** (STRING): Formatted information of disk object that has completed its zeroing operation  
**duration** (INT): The amount of time the zeroing operation required, in seconds  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.disk.zero.fake**

### **Severity**

NOTICE

### **Description**

This message occurs when disk zeroing has been faked on a disk.

### **Corrective Action**

(Call support).



## Syslog Message

%s: fake disk zeroing complete

## Parameters

**disk\_info** (STRING): Formatted information of the disk on which zeroing has been faked

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.diskadd events

## raid.diskadd.abort

### Severity

ERROR

### Description

This message occurs when the system aborts the addition of disks to an aggregate.

### Corrective Action

Use the reason provided to determine an action from the following list: - Change in raidgroup size: Retry adding disks to the volume. - One of the target plexes has failed: Delete the failed plex. Use the "storage aggregate plex delete" command. When the failed plex is removed and the aggregate state appears as "unmirrored", retry adding disks to the volume. - For further information or assistance, contact NetApp technical support.

## Syslog Message

Addition of disks to %s %s%s aborted %s.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Reason the disk addition was aborted.

## raid.diskadd.add.abort

### Severity

ERROR

### Description

This message occurs when the addition of disks to an aggregate is aborted due to a detected disk failure.

## Corrective Action

Retry adding disks to the aggregate. If disk list was specified for addition, replace the failed disk with a good disk in the list.

## Syslog Message

Addition of disks to %s %s%s aborted due to failure of %s.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): String representing the name of the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.diskadd.create.abort

### Severity

ERROR

### Description

This message occurs when the creation of a an aggregate is aborted due to a detected disk failure.

### Corrective Action

Retry the aggregate creation. If disk list was specified for creation, replace the failed disk with a good disk in the list.

### Syslog Message

Creation of %s %s%s aborted due to failure of %s.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the disk shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): String representing the name of the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.disktoc events

### raid.disktoc.rawsizeChange

#### Severity

INFORMATIONAL

#### Description

This event is generated when it is detected from the disk Table Of Contents that a disk's stored rawsize does not match the physical size of the disk. (see burt 81939 for more details).

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**disk\_info** (STRING): The Formatted information of the disk

**label\_rawsize** (LONGINT): The rawsize according to the label.

**phys\_rawsize** (LONGINT): The rawsize according to the storage layer

**instanceFile** (STRING): The internal file code

**instanceld** (INT): The internal instance identifier

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

### raid.disktoc.tooRecent

#### Severity

ALERT

#### Description

This message occurs when RAID discovers that the disk's table of contents version is newer than that supported by this release of Data ONTAP®. Data ONTAP takes appropriate recovery actions, as described in additional logged events.

## Corrective Action

If this message was seen during boot, reboot with the previous kernel version. If this message was seen when the system was up and running, it is likely due to adding a disk to a running system. Remove the disk from system. If the intent is to add the disk as a spare disk, use the 'priv set advanced;unfail -s disk\_name' command to mark the disk as not zeroed spare. NetApp technical support For information about correcting the problem, search for the "raid.disktoc.tooRecent" keyword on the Knowledgebase of the NetApp Support Site.

## Syslog Message

%s has a disk table of contents version (%d) more recent than the currently supported version (%d).

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**disktocVersion** (INT): Disk table of contents version.

**disktocCurrentVersion** (INT): Current disk table of contents version supported by the software.

**instanceFile** (STRING): Internal file code.

**instanceId** (INT): Internal instance identifier.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.enable events

## raid.enable.remote.nvlog

### Severity

INFORMATIONAL

### Description

This event is issued when RAID sub system enables recording in the partner's NVRAM entries for the log based file system.

### Corrective Action

Need to replace NVRAM if observed repeatedly during different boot-ups.

### Syslog Message

Enabled Remote NVRAM logging; status: %d.

### Parameters

**status** (INT): Remote NVRAM log enable status

# raid.encrypt events

## raid.encrypt.disabled

### Severity

ALERT

### Description

This message occurs when a node is enabled for encryption by RAID, but the system is not FIPS compliant. Volume encryption and decryption will not be available because the system is not considered secure. Encrypted volumes will not be accessible until this problem is solved.

### Corrective Action

Reboot the system. If the issue persists after a reboot contact NetApp technical support for assistance.

### Syslog Message

RAID encryption is not available because the system is not FIPS compliant. Error: %d.

### Parameters

**error\_code** (INT): Reason for non FIPS compliance.

## raid.encrypt.no.key

### Severity

ERROR

### Description

This message occurs when an encrypted volume is being mounted and the encryption key can not be found. The volume is inaccessible until this problem is solved.

### Corrective Action

Verify that the network is accessible. Verify that the key manager server is available by using the "security key-manager external show-status" command.

### Syslog Message

No encryption key is available for volume '%s' because of error '%d'.

### Parameters

**volume** (STRING): Volume involved.

**error** (INT): Reason code for the failure to find a key.

# raid.fdr events

## raid.fdr.failed.ok

### Severity

INFORMATIONAL

### Description

This message occurs when we detect an entry in the failed disk registry, and have forcibly failed the disk.

## Corrective Action

Please make sure that the event message `raid.disk.unload.done` or `raid.carrier.remove` has been seen in the log. And then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

## Syslog Message

%s successfully deleted from spare pool

## Parameters

**disk\_info** (STRING): Formatted information of the disk object

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.fdr.full

### Severity

ERROR

### Description

This message occurs when we determine that the failed disk registry is full and we must eliminate an existing disk entry.

### Corrective Action

Please make sure that the event message `raid.disk.unload.done` or `raid.carrier.remove` has been seen in the log. And then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

### Syslog Message

Failed disk registry full! Recycling entry for %s model %s (S/N %s)

### Parameters

**vendor** (STRING): The name of the vendor

**model** (STRING): The model number

**serialno** (STRING): The serial number

## raid.fdr.inservice

### Severity

NOTICE

## Description

This message occurs when the system detects an entry in the failed disk registry, but the disk appears to be incorporated into an aggregate.

## Corrective Action

(None).

## Syslog Message

WARNING: %s previously failed, but is currently in service (%s %s).

## Parameters

**disk\_info** (STRING): Formatted information about the disk object.

**vol\_type** (STRING): Volume type, which is always aggregate.

**vol** (STRING): Name of the aggregate.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.fdr.reminder

### Severity

ERROR

### Description

This message occurs as a reminder that a previously detected failed disk is still present in the system and needs to be removed.

### Corrective Action

The failed disk carrier is fully prepared for removal. Remove and replace the failed disk carrier.

### Syslog Message

Failed %s is still present in the system and should be removed.

### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.fdr.update.error

### Severity

NOTICE

### Description

This message occurs when RAID fails to add an entry for a failed disk into the failed disk registry. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Unable to %s to %s (error %d).

### Parameters

**op** (STRING): Type of add operation that failed.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.

**errorCode** (INT): Error code.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.flexvol events

### raid.flexvol.lock.conflict

#### Severity

INFORMATIONAL

#### Description

This message occurs when a configuration operation on an aggregate (for example, adding disks, relocating aggregate disks, etc.), fails to get exclusive lock access because of one or more indirect conflict volume locks held by child volumes within the aggregate.

#### Corrective Action

Retry the failed operation after the current operation has completed.



## Syslog Message

Exclusive access on the aggregate %s by operation %s failed because of an indirect conflict flexvol lock with current operation %s.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**attempted\_op\_name** (STRING): Name of the operation that tried to get exclusive access.

**current\_op\_name** (STRING): Name of the operation that is currently holding exclusive access.

# raid.fm events

## raid.fm.disasterSummary

### Severity

ERROR

### Description

This event contains summary information for the disaster takeover. It is only generated if a disaster takeover occurs.

### Corrective Action

(None).

## Syslog Message

RAID disaster takeover summary: number of partner volumes and aggregates=%d, number of rewrite-fsid partner volumes and aggregates=%d, number of out-of-date partner volumes and aggregates=%d, number of ignored partner volumes and aggregates=%d, number of local volumes and aggregates=%d, number of out-of-date local volumes and aggregates=%d

## Parameters

**nPvols** (INT): The number of partner volumes involved in the takeover

**nRwPvols** (INT): The number of rewrite-fsid partner volumes

**nOodPvols** (INT): The number of out-of-date partner volumes

**nIgnPvols** (INT): The number of ignored partner volumes

**nLvols** (INT): The number of local volumes involved in the takeover

**nOodLvols** (INT): The number of out-of-date local volumes

## raid.fm.fixedMixedOwnership

### Severity

NOTICE

### Description

This message occurs when assimilation detects that an aggregate contains disks with multiple current owner IDs and fixes the ownership issue.

### Corrective Action

(None).

## Syslog Message

Aggregate %s was detected to have disks with mixed current owner IDs and the ownership has been fixed.

## Parameters

**aggregate** (STRING): Name of the aggregate that contained disks with mixed owner IDs.

**disk\_count** (INT): Number of disks that had their current owner ID fixed.

**current\_owner\_id** (INT): NVRAM system ID of the aggregate's current owner.

## raid.fm.leftBehindAggr

### Severity

NOTICE

### Description

This message occurs when an aggregate is not returned to its original owner in a giveback operation.

### Corrective Action

Run the 'storage failover giveback' command to return the aggregate to its original owner.

## Syslog Message

Aggregate %s was left behind during giveback.

## Parameters

**aggregate** (STRING): Name of the aggregate that was left behind.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.

**current\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's current owner.

## raid.fm.replayFail

### Severity

NOTICE

### Description

This message occurs when RAID replay fails for partner aggregates during takeover. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

RAID takeover: RAID replay failed. Status: %d.

## Parameters

**status** (INT): Replay failure status code.

## raid.fm.takeoverFail

### Severity

ALERT

## Description

This message occurs when a high-availability (HA) takeover attempt fails due to an assimilation error.

## Corrective Action

If partner aggregates are failing to assimilate on the surviving node of the HA pair, contact technical support for assistance with bringing these aggregates online.

## Syslog Message

RAID takeover failed: %s.

## Parameters

**reason** (STRING): Reason the takeover failed.

## raid.fm.volDisasterBadState

### Severity

ERROR

### Description

This message occurs when the system cannot split an out-of-date volume due to a bad volume state.

### Corrective Action

(None).

### Syslog Message

Partner volume %s identity unaltered during HA disaster takeover: bad state %s. Clients might see stale data.

### Parameters

**vol** (STRING): Name of the aggregate.

**state** (STRING): Aggregate state.

## raid.fm.volDisasterFail

### Severity

ERROR

### Description

This message occurs when the system encounters an error during a high availability (HA) disaster takeover.

### Corrective Action

(None).

### Syslog Message

Partner volume %s HA disaster takeover error (%s): %s. Clients will not be able to access this volume.

### Parameters

**vol** (STRING): Name of the aggregate.

**type** (STRING): Type of error

**reason** (STRING): Internal reason code

## raid.fm.volDisasterIgnore

### Severity

ERROR

### Description

This message occurs when the system ignores a volume during a high availability (HA) disaster takeover.

### Corrective Action

(None).

### Syslog Message

Partner volume %s identity unaltered during HA takeover: %s. Clients might see stale data.

### Parameters

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Reason for ignoring the volume

**code** (INT): Internal code

## raid.fm.volDisasterWarn

### Severity

ERROR

### Description

This message occurs when the system encounters a nonfatal error during a high availability (HA) disaster takeover

### Corrective Action

(None).

### Syslog Message

Partner volume %s HA disaster takeover warning (%s): %s Volume takeover might be incomplete and clients might see stale data.

### Parameters

**vol** (STRING): Name of the aggregate.

**type** (STRING): Type of error

**reason** (STRING): Internal reason code

## raid.fm.volFsidRewrite

### Severity

ERROR

### Description

This message occurs when the system rewrites a volume's or aggregate's canceling giveback high availability (HA) disaster takeover.

### Corrective Action

Remount the clients of the volume.

## Syslog Message

Partner %s %s FSID was rewritten during HA disaster takeover. Clients must be remounted.

## Parameters

**vol\_type** (STRING): Volume type.

**vol** (STRING): Name of the aggregate.

**oldFsid** (INT): Old fsid

**newFsid** (INT): New fsid

## raid.fm.volFsidRewriteOod

### Severity

ERROR

### Description

This message occurs when the system rewrites a volume's File System ID (FSID) during a high availability (HA) disaster takeover with an out-of-date volume.

### Corrective Action

Remount the clients of the volume.

## Syslog Message

Partner volume %s FSID was rewritten during HA disaster takeover. Clients must be remounted.

## Parameters

**vol** (STRING): Name of the aggregate.

**oldFsid** (INT): Old fsid

**newFsid** (INT): New fsid

## raid.foreignaggregate events

### raid.foreignAggregate.online.rtidChanged

### Severity

NOTICE

### Description

This message occurs when the system changes the RAID tree ID of a foreign aggregate that is being brought online.

### Corrective Action

(None).

## Syslog Message

%s: RAID tree ID was changed from %s to %s.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**oldrtid** (STRING): Old RAID tree ID of the aggregate.

**newrtid** (STRING): New RAID tree ID of the aggregate.

# raid.fs events

## raid.fs.disk.ownership.change

### Severity

NOTICE

### Description

This message occurs when a filesystem disk is prepared for ownership change by administrator forcibly.

### Corrective Action

(None).

### Syslog Message

Ownership of %s, which is part of an aggregate %s, is being changed by administrator forcibly.

### Parameters

**disk\_info** (STRING): Formatted information of the disk.

**vol\_name** (STRING): Name of the aggregate.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.hybrid events

## raid.hybrid.SSDHAExceed

### Severity

ERROR

### Description

This message occurs when the sum of all sizes of solid state disks (SSDs) disks of hybrid aggregates on this node and its partner node is larger than the maximum system limit. Takeover might fail because of this. It is issued by the Storage Health Monitor.

### Corrective Action

Bring the sum of all sizes of SSDs of hybrid aggregates on this node and its partner below the limit. If you cannot bring the sum of all sizes of SSDs of hybrid aggregates on both nodes below the limit, disable HA failover by using the 'storage failover modify -node [nodename] -enabled false' command. However, this compromises the high availability of your system and is only recommended for use until the size issue has been resolved. To avoid this problem in the future, be careful when adding SSDs to existing hybrid aggregates not to exceed the limit for both nodes in the HA pair.

## Syslog Message

The sum of all sizes of SSD disks of hybrid aggregates on this node and its partner %s, exceeds %s. Takeover might fail because of this.

## Parameters

**hybrid\_ssd\_disks\_sizes** (STRING): Sum of sizes of SSDs disks of hybrid aggregates on this node and its partner node.

**limit** (STRING): Maximum system limit total of sizes of SSDs of hybrid aggregates.

## raid.hybrid.SSDTotExceed

### Severity

ERROR

### Description

This message occurs as a reminder that the sum of all sizes of SSD disks of hybrid aggregates is too large. It is issued by Storage Health Monitor.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

The sum of sizes of SSD disks of hybrid aggregates %s exceeds the %s maximum. Contact NetApp technical support.

## Parameters

**ssd\_disks\_size** (STRING): Total of sizes of SSD disks of hybrid aggregates.

**limit** (STRING): Maximum total of sizes of SSD disks of hybrid aggregates.

## raid.initialize events

### raid.initialize.nonhomed.disk

### Severity

NOTICE

### Description

This message occurs when a foreign or partner homed spare disk is detected during system initialization.

### Corrective Action

In Maintenance mode, remove this disk's ownership by using the 'disk remove\_ownership disk\_name' command and assign it to the intended node by using the 'disk assign disk\_name' command.

## Syslog Message

%s: foreign or partner homed spare disk detected during system initialization.

## Parameters

**disk\_info** (STRING): Formatted information of the foreign spare disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.label events

### raid.label.io.readError

#### Severity

NOTICE

#### Description

This message occurs when a RAID label read I/O failure is detected. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

#### Corrective Action

(None).

#### Syslog Message

Label read on %s failed with storage error %s. The system will stop using the disk for I/O operations.

#### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**error\_str** (STRING): Error string.

**instanceFile** (STRING): Internal file code.

**instanceId** (INT): Internal instance identifier.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Character string representing the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

### raid.label.io.revertError

#### Severity

NOTICE

#### Description

This message occurs when a RAID label revert to consistent state fails during assimilation. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.



## Corrective Action

(None).

## Syslog Message

Label revert on plex %s, %s failed.

## Parameters

**plex** (STRING): Name of the containing plex.

**disk\_info** (STRING): Formatted information about the disk.

**error\_str** (STRING): Error string.

**instanceFile** (STRING): Internal file code.

**instanceId** (INT): Internal instance identifier.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Character string representing the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.label.io.writeError

### Severity

NOTICE

### Description

This message occurs when a RAID label write I/O failure is detected. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Label write on %s failed with storage error %s. The system will stop using the disk for I/O operations.

### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**error\_str** (STRING): Error string.

**instanceFile** (STRING): Internal file code.

**instanceId** (INT): Internal instance identifier.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Character string representing the model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.label.raidarea.invalidVersion

### Severity

NOTICE

### Description

This message occurs when the system discovers that a disk has a RAIDArea object version of 0, which is invalid. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s has %s RAIDArea object with a version (%d) that is not supported (the current object version is %d).

### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**type** (STRING): RAIDArea object type.

**version** (INT): RAIDArea object version.

**currentVersion** (INT): Current version of this RAIDArea object type supported by the software.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Character string representing the model of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.label.raidarea.tooRecent

### Severity

NOTICE

### Description

This message occurs when the system discovers that a disk has an invalid RAIDArea object version. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s has %s raidarea object with a version (%d) more recent than the currently supported version (%d).

## Parameters

**disk\_info** (STRING): Formatted information about the disk.  
**type** (STRING): RAIDArea object type.  
**version** (INT): RAIDArea object version.  
**currentVersion** (INT): Current version of this RAIDArea object type supported by the software.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Character string representing the model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.labeledit events

### raid.labeledit.buildtree

#### Severity

NOTICE

#### Description

This event occurs when you run the 'label buildtree' command to create an aggregate from a collection of disks.

#### Corrective Action

(None).

#### Syslog Message

%s %s was created using the 'label buildtree' command.

#### Parameters

**voltype** (STRING): Volume type.  
**volname** (STRING): Name of the aggregate created by the command.  
**cmdline** (STRING): Command line used to create the aggregate.

## raid.labelmaint events

### raid.labelmaint.isolate.aggr

#### Severity

INFORMATIONAL

#### Description

This message occurs when the system places an aggregate into online label maintenance mode.

**Corrective Action**

(None).

**Syslog Message**

Placing volume %s%s into online label maintenance mode.

**Parameters**

**owner** (STRING): Owner of the disk drive.

**vol** (STRING): Name of the aggregate.

**raid.labelmaint.isolate.disk****Severity**

INFORMATIONAL

**Description**

This message occurs when the system places a disk drive into online label maintenance mode. Online label maintenance is a utility available in normal mode that enables users to perform operations on disk drives that typically are allowed only in maintenance mode. Examples of such operations are label editing and volume recovery.

**Corrective Action**

(None).

**Syslog Message**

Placing %s%s into online label maintenance mode.

**Parameters**

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk drive.

**raid.labelmaint.unisolate****Severity**

INFORMATIONAL

**Description**

This message occurs when the system takes a disk drive out of online label maintenance mode.

**Corrective Action**

(None).

**Syslog Message**

Taking %s%s out of online label maintenance mode.

**Parameters**

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk drive.

# raid.mcc events

## raid.mcc.disasterSummary

### Severity

INFORMATIONAL

### Description

This message reports the summary information for the MetroCluster(tm) switchover operation.

### Corrective Action

(None).

### Syslog Message

RAID switchover summary: number of DR aggregates=%d number of mirrored aggregates=%d number of unmirrored aggregates=%d number of out-of-date aggregates=%d number of mirrored aggregates with both plexes online=%d number of offline aggregates switched over=%d

### Parameters

**nTotalAggrs** (INT): Number of DR aggregates switched over.

**nMirAggrs** (INT): Number of mirrored aggregates switched over.

**nUnmirAggrs** (INT): Number of unmirrored aggregates switched over.

**nOodPlexes** (INT): Number of mirrored aggregates with out-of-date plexes.

**nMirComplete** (INT): Number of mirrored aggregates with both the plexes online during switchover.

**nOfflineAggrs** (INT): Number of offline aggregates switched over.

## raid.mcc.heal.aggrs.canceled

### Severity

ERROR

### Description

This message occurs when a 'metrocluster heal -phase aggregates' operation is canceled due to the presence of left-behind high-availability (HA) aggregates.

### Corrective Action

Giveback still in progress can cause this condition. To see if this is the case, execute 'storage failover show-giveback'. If it is in progress, wait until it completes and retry the 'metrocluster heal -phase aggregates' command. If giveback is not in progress then execute a 'storage failover giveback -ofnode' from the node currently hosting the left-behind aggregate. After this is complete, retry the 'metrocluster heal -phase aggregates' command.

### Syslog Message

Aggregate %s homed on %s is a %s, canceling heal aggregates operation.

### Parameters

**aggregate** (STRING): Name of the aggregate

**home\_owner\_name** (STRING): Name of the node home to the left behind aggregate.

**reason** (STRING): Activity that is preventing heal aggregates operation.

## raid.mcc.heal.roots.canceled

### Severity

ERROR

### Description

This message occurs when a 'metrocluster heal -phase root-aggregates' operation is canceled because a switched over root aggregate is busy with another configuration operation. An ongoing configuration operation can cause this condition.

### Corrective Action

Wait a few minutes, then retry the 'metrocluster heal -phase root-aggregates' command.

### Syslog Message

Aggregate %s is busy with another configuration operation. Canceling heal root-aggregates operation.

### Parameters

**aggregate** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the heal root-aggregates operation.

## raid.mcc.ISL.down

### Severity

ERROR

### Description

This message occurs when a MetroCluster(tm) DR operation other than disaster switchover is canceled because remote site storage is not visible; this can happen if the Inter-Switch Links are down.

### Corrective Action

1. Verify remote cluster connectivity. Use the "storage switch add -switch-ipaddress <switch-ip-address>" command to start monitoring the switches if not already so. Then use the "storage switch show -port" and "storage switch show -san-config" commands to verify FC connectivity to the remote site. 2. Retry the operation. If the error persists, contact technical support for assistance.

### Syslog Message

%s was canceled because remote site storage is not visible; Inter-Switch Links might be down.

### Parameters

**operation** (STRING): The DR operation which was canceled.

## raid.mcc.leftBehindAggr

### Severity

NOTICE

### Description

This message occurs when an aggregate is not returned to its original owner in a switchback operation.

### Corrective Action

Run the 'metrocluster switchback' command to return the aggregate to its original owner. If the condition persists, contact NetApp technical support.

## Syslog Message

Aggregate %s was left behind during switchback.

## Parameters

**aggregate** (STRING): Name of the aggregate that was left behind.  
**aggregate\_uuid** (STRING): Universal Unique Identifier(UUID) of the aggregate.  
**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.  
**current\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's current owner.  
**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's disaster recovery(DR) home owner.

## raid.mcc.leftBehindSpare

### Severity

NOTICE

### Description

This message occurs when a spare disk is not returned to its original owner in a switchback operation.

### Corrective Action

Run the 'metrocluster switchback' command to return the disk to its original owner.

## Syslog Message

Disk %s was left behind during switchback.

## Parameters

**disk** (STRING): Name of the disk that was left behind.  
**disk\_uuid** (STRING): Universal Unique Identifier(UUID) of the disk.  
**home\_owner\_id** (LONGINT): NVRAM system ID of the disk's home owner.  
**current\_owner\_id** (LONGINT): NVRAM system ID of the disk's current owner.  
**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the disk's disaster recovery(DR) home owner.

## raid.mcc.nso.canceled

### Severity

ERROR

### Description

This message occurs when a MetroCluster negotiated switchover is canceled due to activity in RAID on a volume or aggregate.

### Corrective Action

Reissue the 'metrocluster switchover' command after the pending or in-progress operation on the aggregate is complete. The 'override-vetoes' option may be used to proceed without waiting for the pending or in-progress operation to complete.

## Syslog Message

Aggregate %s is %s, canceling negotiated switchover.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing negotiated switchover.

## **raid.mcc.nso.dr.leftbehind**

### **Severity**

ERROR

### **Description**

This message occurs when a MetroCluster switchover is canceled due to the presence of left behind aggregates from a switchback operation.

### **Corrective Action**

If the aggregate is not in a "failed" or "limbo" state, issue the 'metrocluster switchback' command to return the aggregate to its original owner. Reissue the 'metrocluster switchover' command.

### **Syslog Message**

Aggregate %s was left behind from a switchback operation, canceling negotiated switchover.

### **Parameters**

**aggregate** (STRING): Name of the aggregate that was left behind.

**aggregate\_uuid** (STRING): Universal Unique Identifier of the aggregate.

**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.

**current\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's current owner.

**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's disaster recovery home owner.

## **raid.mcc.nso.ha.leftbehind**

### **Severity**

ERROR

### **Description**

This message occurs when a MetroCluster switchover is canceled due to the presence of left behind aggregates from a giveback operation.

### **Corrective Action**

If the aggregate is not in a "failed" or "limbo" state, issue the 'storage failover giveback' command to return the aggregate to its original owner. Reissue the 'metrocluster switchover' command.

### **Syslog Message**

Aggregate %s was left behind from a giveback operation, canceling negotiated switchover.

### **Parameters**

**aggregate** (STRING): Name of the aggregate that was left behind.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.

**current\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's current owner.

## **raid.mcc.remote.storage.down**

### **Severity**

ERROR



## Description

This message occurs when a MetroCluster(tm) DR operation other than disaster switchover is canceled because remote site storage is not visible; this can happen if the FC or SAS links to the remote storage are down.

## Corrective Action

1. Verify remote cluster connectivity. Use the "fcp interface show" command to verify FC connectivity to the remote site.
2. Retry the operation. If the error persists, contact technical support for assistance.

## Syslog Message

%s was canceled because remote site storage is not visible; FC or SAS links might be down.

## Parameters

**operation** (STRING): The DR operation which was canceled.

## raid.mcc.root.configError

### Severity

ALERT

### Description

This message occurs when the system detects a configuration error on the root aggregate in a MetroCluster(tm) configuration. Remote disks and local disks are not being used in the proper configuration. One plex should comprise only disks from the local pool, Pool0, and the other plex should comprise only disks from the remote pool, Pool1. If the inter-site communication fails between the MetroCluster sites, access to remote disks will be lost. This creates an elevated risk of local site downtime.

### Corrective Action

Replace all remote (Pool1) disks in the local plex with local (Pool0) spares. In the remote plex, replace all local (Pool0) disks with remote (Pool1) spares. Use the 'storage disk replace' command to initiate disk replacements.

### Syslog Message

In the root aggregate, %s

### Parameters

**error\_message** (STRING): Description of which plexes have disks from the wrong pool.

## raid.mcc.root.unmirrored

### Severity

ALERT

### Description

This message occurs when the root aggregate is not mirrored in a MetroCluster(tm) configuration.

### Corrective Action

Mirror the root aggregate by using the 'storage aggregate mirror' command.

### Syslog Message

The root aggregate is not mirrored. This aggregate must be mirrored in a MetroCluster configuration or a switchover of this site might not be possible in case of a disaster.

## Parameters

(None).

## raid.mcc.switchbackCanceled

### Severity

ERROR

### Description

This message occurs when a MetroCluster switchback is canceled due to activity in RAID on a disaster partner volume or aggregate.

### Corrective Action

Reissue the 'metrocluster switchback' command after the pending or in-progress operation on the DR partner aggregate is complete. The override-vetoes option may be used to proceed without waiting for the pending or in-progress operation to complete.

### Syslog Message

DR aggregate %s is %s, canceling switchback.

### Parameters

**vol** (STRING): Name of the aggregate

**reason** (STRING): Activity that is preventing DR switchback

## raid.mcc.switchbackSummary

### Severity

INFORMATIONAL

### Description

This message occurs when a node has finished releasing switched over aggregates and spare disks during switchback.

### Corrective Action

(None).

### Syslog Message

RAID switchback summary: switched back aggregates: %d, left behind aggregates: %d, switched back spares: %d

### Parameters

**switched\_back\_aggr\_count** (INT): Count of aggregates that were switched back.

**left\_behind\_aggr\_count** (INT): Count of switched over aggregates that remain after switchback.

**switched\_back\_spare\_count** (INT): Count of spare disks that were switched back.

## raid.mgrrio events

### raid.mgrrio.lw.id.inconsist

## Severity

NOTICE

## Description

This message occurs during a SyncMirror® resynchronization operation, when the system detects an inconsistent RAID write signature on a RAID stripe. This RAID inconsistent write signature might be due to a write operation on disk that was not committed successfully. Data ONTAP® attempts to recover such a bad block automatically. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

## Corrective Action

(None).

## Syslog Message

Inconsistent RAID write signature on RAID group %s%s, stripe #%llu, during SyncMirror resync operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

# raid.mgvio events

## raid.mgvio.lw.id.inconsist

## Severity

NOTICE

## Description

This message is generated after detecting an inconsistent RAID write signature on a RAID stripe, during SyncMirror verify operation. This inconsistent RAID write signature may be due to a write operation on disk was not committed successfully. Data ONTAP will attempt to recover such bad block automatically. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

## Corrective Action

(None).

## Syslog Message

Inconsistent RAID write signature on RAID group %s%s, stripe #%llu, during a SyncMirror verify operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

# raid.mirror events

## raid.mirror.aggrSnapUse

### Severity

NOTICE

### Description

This message occurs when a SyncMirror® aggregate uses aggregate Snapshot(tm) copies other than SyncMirror aggregate Snapshot copies. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after one plex was temporarily offline. However, Data ONTAP® sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another. To reduce the chance of deleting SyncMirror Snapshot copies, you should not create or schedule aggregate Snapshot copies in SyncMirror aggregates.

### Corrective Action

Disable automatic aggregate Snapshot copies and delete all Snapshot copies in the aggregate except those with names starting with "mirror\_resync" or "mirror\_verify". To disable automatic aggregate Snapshot copies, use the 'snap sched -A aggr-name 0' command. To list Snapshot copies in the aggregate, use the 'snap list -A aggr-name' command. To delete a Snapshot copy, use the 'snap delete -A aggr-name snapshot-name' command. For further information or assistance, contact NetApp technical support.

### Syslog Message

Aggregate Snapshot copies are used in SyncMirror %s '%s%s'. Creating or scheduling Snapshot copies in SyncMirror aggregates is not recommended.

### Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the warning is issued about an aggregate.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.mirror.bigio.restrict

### Severity

ALERT

### Description

An aggregate that experienced a medium error during reconstruction is restricted and marked wafli-inconsistent, but starting wafliion has failed. This event is issued to alert operator that aggregate is not accessible and wafliion must be started to allow access to it.

### Corrective Action

Start wafliion on the indicated aggregate.

### Syslog Message

%s %s is restricted. Start wafliion to allow access to it

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## raid.mirror.bigio.restrict.failed

### Severity

EMERGENCY

### Description

This message occurs when automatically restricting an aggregate fails after a medium error happened during a reconstruction.

### Corrective Action

Contact NetApp technical support. For further information about correcting the problem, see the knowledgebase article 3013638.

### Syslog Message

Failed to restrict the %s %s (%s)

### Parameters

**vol\_type** (STRING): Volume type.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Specific reason preventing the restrict operation.

## raid.mirror.bigio.wafiron.nostart

### Severity

EMERGENCY

### Description

This message occurs when wafiron fails to start automatically on an aggregate that experienced a medium error during reconstruction and is restricted and marked WAFL® inconsistent.

### Corrective Action

Contact NetApp technical support. For further information about correcting the problem, see the knowledgebase article 3013638.

### Syslog Message

Wafiron cannot start on %s %s (%s)

### Parameters

**vol\_type** (STRING): Volume type.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Specific reason preventing wafiron to start.

## raid.mirror.faultIsolation.reminder

### Severity

ERROR

### Description

This message occurs when both plexes of a SyncMirror® configuration have disks in the same hardware-based disk pool.

## Corrective Action

a) Identify the problem aggregate from the system logs. b) Determine which mirror plexes have disks in the same hardware-based pools. c) Determine how this occurred: possible causes include 1) Wiring problem 2) Reconstruction forced mixed pools 3) Mirror was created forcibly d) Based on the information you gather, determine how to correct the issue. For example, use the "storage disk replace" command to copy the disk belonging to the wrong pool to a disk belonging to the right pool. e) If you need assistance, contact NetApp technical support.

## Syslog Message

%s %s plexes not fault isolated. Multiple plexes have disks in: %s

## Parameters

**voltype** (STRING): Aggregate or volume.

**volname** (STRING): Name of the aggregate or volume.

**pool** (STRING): Disk pools that have disks from the aggregate.

## raid.mirror.lowSnapReserve

### Severity

ERROR

### Description

This message occurs when the aggregate Snapshot(tm) copy reserve in a SyncMirror® aggregate is too low, increasing the risk of deleting SyncMirror Snapshot copies. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after a temporary loss of connectivity to one plex. However, Data ONTAP® sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another.

### Corrective Action

Increase the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'. Do not decrease the aggregate Snapshot copy reserve in SyncMirror aggregates below the default 5%. The suggested Snapshot copy reserve might vary from message to message for the same aggregate if the write load on the aggregate changes and especially if the aggregate option 'resyncsnaptime' is set to a significantly lower value than the default 60 minutes. If one suggestion seems unreasonably high, you might want to track messages for several days and set the aggregate Snapshot copy reserve to the highest value that is consistently suggested during periods of high write load on the aggregate.

### Syslog Message

Aggregate Snapshot copy reserve in SyncMirror %s '%s%s' is too low. It is set to %d%%. Increase it to %d%%.

### Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the warning is issued about an aggregate.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**current\_snap\_reserve** (INT): Current aggregate Snapshot copy reserve (percent).

**suggested\_snap\_reserve** (INT): Suggested higher aggregate Snapshot copy reserve (percent).

## raid.mirror.read.mismatch

**Severity**

NOTICE

**Description**

This message occurs when the system detects a mismatch between data in two plexes of a SyncMirror® aggregate. The system does not perform data verification by reading from both plexes during normal operation, but it does so during wafiron. The system fixes the inconsistency across mirrored plexes.

**Corrective Action**

(None).

**Syslog Message**

Mirror read verification failed in the %s '%s%s': mismatch between disks %s and %s (vbn %llu, blockNum %llu).

**Parameters**

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**srcDisk** (STRING): Name of the disk in the first plex.  
**dstDisk** (STRING): Name of the disk in the second plex.  
**volumeBno** (LONGINT): Volume block number.  
**blockNum** (LONGINT): Disk block number.

**raid.mirror.readerr.block.rewrite****Severity**

NOTICE

**Description**

This event is issued when mirror read error handling fixes a multi-disk media or checksum error on Raid0 volumes.

**Corrective Action**

(None).

**Syslog Message**

Mirror error handler rewriting bad block on %s%s, block #%llu

**Parameters**

**owner** (STRING): Owner of the affected aggregate.  
**disk** (STRING): The name of the disk containing the block being rewritten.  
**blockNum** (LONGINT): The physical block number of the disks being rewritten.

**raid.mirror.resync.deferred****Severity**

NOTICE

**Description**

This message occurs when a mirror plex resynchronization is deferred due to inadequate incore resources.

**Corrective Action**

(None).

**Syslog Message**

%s%s: resynchronization deferred (%s)

**Parameters**

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): Name of the mirror object that could not resynchronize.

**reason** (STRING): Text describing the reason.

**raid.mirror.resync.deferred.ok****Severity**

NOTICE

**Description**

This message occurs when a previously deferred resynchronization is now ready to proceed because of availability of incore resources.

**Corrective Action**

(None).

**Syslog Message**

%s%s: resynchronization previously deferred, now proceeding

**Parameters**

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): Name of the mirror object that can now resynchronize.

**raid.mirror.resync.done****Severity**

NOTICE

**Description**

This event is issued when resynchronization has been completed on a specific mirror.

**Corrective Action**

(None).

**Syslog Message**

%s: resynchronization completed in %s

**Parameters**

**mirror** (STRING): The name of the mirror object that is completed resynchronization

**duration** (STRING): The amount of time the resynchronization required

**raid.mirror.resync.progress**



**Severity**

NOTICE

**Description**

This message occurs every five minutes while RAID mirror resynchronization is in progress. It also occurs at the end of RAID group synchronization.

**Corrective Action**

(None).

**Syslog Message**

%s: mirror resync is %s%% completed. At current throughput rate of %d MB/s it is estimated to finish in %s (HH:MM:SS).

**Parameters**

**tree** (STRING): RAID mirrored tree name.

**percent** (STRING): Percentage of mirror resync completed.

**throughput** (INT): Current input/output (I/O) rate.

**eta\_duration** (STRING): The amount of time the resynchronization is expected to run.

**raid.mirror.resync.snap.base****Severity**

INFORMATIONAL

**Description**

This message occurs at the beginning of a resynchronization and it lists the base and the resynchronization Snapshot(tm) copies for the resynchronization.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s: base Snapshot %d, CP %d (%d), resync Snapshot %d, CP %d (%d)

**Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**base\_snapid** (INT): Base Snapshot identifier.

**base\_CP\_count** (INT): Base Snapshot consistency point count.

**base\_timestamp** (INT): Base Snapshot timestamp (seconds since 1/1/1970).

**next\_snapid** (INT): Resynchronization Snapshot identifier.

**next\_CP\_count** (INT): Resynchronization Snapshot consistency point count.

**next\_timestamp** (INT): Resynchronization Snapshot timestamp (seconds since 1/1/1970).

**raid.mirror.resync.snapcrtfail****Severity**

ERROR

## Description

This message occurs when resynchronization fails to create a mirror resync Snapshot(tm) copy.

## Corrective Action

Increase the size of the root volume or delete old Snapshot copies.

## Syslog Message

%s %s%s: could not create mirror resynchronization Snapshot copy %s (%s)

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being resynchronized.

**snapName** (STRING): Name of the Snapshot copy that was not created.

**error** (STRING): Error code returned by the failed operation.

## raid.mirror.resync.snapdelfail

### Severity

NOTICE

### Description

This message occurs when resynchronization fails to delete a mirror resync Snapshot(tm)copy. The WAFL® Snapshot copy autodelete functionality will automatically delete the Snapshot copy.

### Corrective Action

(None).

### Syslog Message

%s %s%s: could not delete mirror resynchronization Snapshot copy %s (%s)

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being resynchronized.

**snapName** (STRING): Name of the Snapshot copy that was not deleted.

**error** (STRING): Error code returned by the failed operation.

## raid.mirror.resync.snaprenamefail

### Severity

NOTICE

### Description

This message occurs when resynchronization has failed to rename a mirror resync Snapshot(tm) copy. The Snapshot copy is marked as invalid and is deleted.

### Corrective Action

(None).

## Syslog Message

%s %s%s: could not rename mirror resynchronization Snapshot copy %s to %s (%s)

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being resynchronized.

**snapName** (STRING): Name of the Snapshot copy that was not renamed.

**snapName2** (STRING): Attempted new name of the Snapshot copy.

**error** (STRING): Error code returned by the failed operation.

## raid.mirror.resync.snaprenameok

### Severity

INFORMATIONAL

### Description

This message occurs when resynchronization renames a mirror resync Snapshot(tm) copy.

### Corrective Action

(None).

## Syslog Message

%s %s%s: renamed mirror resynchronization Snapshot copy %s to %s

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being resynchronized.

**snapName** (STRING): Old name of the Snapshot copy that has been renamed.

**snapName2** (STRING): New name of the Snapshot copy that has been renamed.

## raid.mirror.resync.start

### Severity

NOTICE

### Description

This message occurs when resynchronization is initiated on a specific mirror.

### Corrective Action

(None).

## Syslog Message

%s%s: start resynchronize to target %s

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): Name of the mirror object that is initiating resynchronization.

**plex** (STRING): Name of the plex object that is the resynchronization target.

## raid.mirror.snapDel.degraded

### Severity

ALERT

### Description

This message occurs when a SyncMirror® aggregate Snapshot(tm) copy is deleted while one plex is offline or resynchronizing. Fast resynchronization of that plex (level 1) is no longer possible. The plex must be resynchronized by copying all data from the online plex (level 0). Without an intervention, it is likely that the level 0 resynchronization will also fail repeatedly, and that it will never be completed. SyncMirror uses aggregate Snapshot copies for resynchronization after a temporary loss of connectivity to one plex. Data ONTAP® sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. This message indicates that configuration of aggregate space use must be changed for SyncMirror resynchronization to be completed.

### Corrective Action

1. Disable automatic aggregate Snapshot copies in the aggregate using the command 'snap sched -A aggr-name 0'. 2. Get a list of Snapshot copies in the aggregate using the command 'snap list -A aggr-name'. Delete all Snapshot copies in the aggregate except those with names starting with "mirror\_resync". 3. Increase the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'. Increase the aggregate Snapshot copy reserve by as much free space as the aggregate allows, but there is no need to increase it beyond 50%. 4. If further increasing the aggregate Snapshot copy reserve is not possible, disable automatic deletion of aggregate Snapshot copies using the command 'aggr options snapshot\_autodelete off'. That might also disable space guarantees on flexible volumes in the aggregate. In that case, you must monitor space used in the aggregate with the command 'df -A aggr-name'. If the aggregate gets full, the application write operations will fail. In environments that are sensitive to that error, such as CIFS or LUNs, disabling automatic deletion of aggregate Snapshot copies should be avoided if possible. If you enable automatic deletion of aggregate Snapshot copies using the command 'aggr options snapshot\_autodelete on', the plex resynchronization will probably fail again. 5. If resynchronization cannot be completed without filling up the aggregate, consider adding more disks to the aggregate. 6. After successful resynchronization, enable automatic deletion of aggregate Snapshot copies using the command 'aggr options snapshot\_autodelete on'. Also restore the original aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'. Do not decrease the aggregate Snapshot copy reserve in SyncMirror aggregates below the default 5%.

### Syslog Message

The SyncMirror aggregate Snapshot copy in %s '%s%s' is being deleted while the aggregate is mirror-degraded. Level 1 resync is not possible.

### Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the message is issued about an aggregate.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.mirror.snapDel.normal

### Severity

ERROR

### Description

This message occurs when a SyncMirror® aggregate Snapshot(tm) copy is deleted while plexes are synchronized. This creates no immediate problem because Data ONTAP® creates a new SyncMirror

aggregate Snapshot copy, but it indicates that the same event can occur when one plex is offline or resynchronizing, which is a problem. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after a temporary loss of connectivity to one plex. Data ONTAP sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another.

### Corrective Action

Increase the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'. Do not decrease the aggregate Snapshot copy reserve in SyncMirror aggregates below the default 5%. The suggested Snapshot copy reserve might vary from message to message for the same aggregate if the write load on the aggregate changes. If one suggestion seems unreasonably high, you might want to track messages for several days and set the aggregate Snapshot copy reserve to the highest value that is consistently suggested during periods of high write load on the aggregate.

### Syslog Message

The SyncMirror aggregate Snapshot copy in %s '%s%s' is being deleted. The aggregate Snapshot copy reserve is set to %d%%. Increase it to %d%%.

### Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the warning is issued about an aggregate.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**current\_snap\_reserve** (INT): Current aggregate Snapshot copy reserve (percent).

**suggested\_snap\_reserve** (INT): Suggested higher aggregate Snapshot copy reserve (percent).

## raid.mirror.snapEst.degraded

### Severity

ERROR

### Description

This message occurs periodically when one plex in a SyncMirror® aggregate is offline or failed, and automatic deletion of aggregate Snapshot(tm) copies is enabled. It provides an estimate of the time before a SyncMirror Snapshot copy might get deleted. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after a temporary loss of connectivity to one plex. However, Data ONTAP® sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another.

### Corrective Action

1. Disable automatic aggregate Snapshot copies in the aggregate using the command 'snap sched -A aggr-name 0'.
2. Get a list of Snapshot copies in the aggregate using the command 'snap list -A aggr-name'. Delete all Snapshot copies in the aggregate except those with names starting with "mirror\_resync".
3. Bring online the plex that is offline or failed as soon as possible.
4. Do not rely only on the estimate of the time before automatic deletion of aggregate Snapshot copies. Monitor space used by the SyncMirror Snapshot copies with the command 'df -A aggr-name'. If SyncMirror Snapshot copies grow beyond the aggregate Snapshot copy reserve, they might be automatically deleted, and that prevents fast SyncMirror resynchronization.

### Syslog Message

SyncMirror %s '%s%s' is mirror-degraded. %d%% of the Snapshot copy reserve is used. SyncMirror Snapshot copy is estimated to be automatically deleted in %s.

## Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the message is issued about an aggregate.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**snap\_reserve\_used** (INT): Percent of Snapshot copy reserve space used.

**est\_time\_left** (STRING): Estimated time until automatic deletion of aggregate Snapshot copies.

## raid.mirror.snapResExpand.failed

### Severity

ERROR

### Description

This message occurs when one plex in a SyncMirror® aggregate is offline, failed, or resynchronizing, and Data ONTAP® attempts to increase the aggregate Snapshot(tm) copy reserve in that aggregate to delay deleting SyncMirror Snapshot copies, but changing the aggregate Snapshot copy reserve fails. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after a temporary loss of connectivity to one plex. However, Data ONTAP sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another.

### Corrective Action

1. Try to increase aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name new-snap-reserve'. 2. Disable automatic aggregate Snapshot copies in the aggregate using the command 'snap sched -A aggr-name 0'. 3. Get a list of Snapshot copies in the aggregate using the command 'snap list -A aggr-name'. Delete all Snapshot copies in the aggregate except those with names starting with "mirror\_resync". 4. Bring online the plex that is offline or failed as soon as possible. 5. Monitor space used by the SyncMirror Snapshot copies with the command 'df -A aggr-name'. If SyncMirror Snapshot copies grow beyond the aggregate Snapshot copy reserve, they might be automatically deleted, and that prevents fast SyncMirror resynchronization. 6. When resynchronization is complete, restore the old value for the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name old-snap-reserve'.

### Syslog Message

An attempt to increase the aggregate Snapshot copy reserve in SyncMirror %s '%s%s' from %d%% to %d%% failed (%s).

## Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the message is issued about an aggregate.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**old\_snap\_reserve** (INT): Old aggregate Snapshot copy reserve (percent).

**new\_snap\_reserve** (INT): New higher aggregate Snapshot copy reserve (percent).

**reason** (STRING): Reason for the failure.

## raid.mirror.snapResExpanded

### Severity

NOTICE

## Description

This message occurs when one plex in a SyncMirror® aggregate is offline, failed, or resynchronizing, and the aggregate Snapshot(tm) copy reserve in that aggregate is increased to delay deleting SyncMirror Snapshot copies. The aggregate Snapshot copy reserve will be reverted to its previous value when resynchronization is complete. You can always change the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'. SyncMirror uses aggregate Snapshot copies to allow fast resynchronization after a temporary loss of connectivity to one plex. However, Data ONTAP® sometimes must delete aggregate Snapshot copies to preserve space guarantees for flexible volumes. If that happens during SyncMirror resynchronization or while one plex is offline, the mirror must be reinitialized by copying all data from one plex to another.

## Corrective Action

1. Disable automatic aggregate Snapshot copies in the aggregate using the command 'snap sched -A aggr-name 0'. 2. Get a list of Snapshot copies in the aggregate using the command 'snap list -A aggr-name'. Delete all Snapshot copies in the aggregate except those with names starting with "mirror\_resync". 3. Bring online the plex that is offline or failed as soon as possible. 4. Monitor space used by the SyncMirror Snapshot copies with the command 'df -A aggr-name'. If SyncMirror Snapshot copies grow beyond the aggregate Snapshot copy reserve, they might be automatically deleted, and that prevents fast SyncMirror resynchronization.

## Syslog Message

Aggregate Snapshot copy reserve in SyncMirror %s '%s%s' is increased from %d%% to %d%% while the mirror is degraded or resyncing. Aggregate Snapshot copy reserve will be reverted to the old value when resync is complete.

## Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the message is issued about an aggregate.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**old\_snap\_reserve** (INT): Old aggregate Snapshot copy reserve (percent).  
**new\_snap\_reserve** (INT): New higher aggregate Snapshot copy reserve (percent).

## raid.mirror.snapResReverted

### Severity

NOTICE

### Description

This message occurs when the aggregate Snapshot(tm) copy reserve in a SyncMirror® aggregate is reverted to the original value after a successful resynchronization in that aggregate.

### Corrective Action

Verify that the new Snapshot copy reserve is correct. You can always change the aggregate Snapshot copy reserve using the command 'snap reserve -A aggr-name percent'.

### Syslog Message

Aggregate Snapshot copy reserve in SyncMirror %s '%s%s' was reverted from %d%% back to %d%%.

### Parameters

**vol\_type** (STRING): Always "aggregate," indicating that the message is issued about an aggregate.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

**old\_snap\_reserve** (INT): Old aggregate Snapshot copy reserve (percent).  
**new\_snap\_reserve** (INT): New aggregate Snapshot copy reserve (percent).

## **raid.mirror.verify.aborted**

### **Severity**

NOTICE

### **Description**

This event is issued when verification has been completed on a specific mirror due to abort.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s: verification stopped after %s

### **Parameters**

**owner** (STRING): Owner of the affected aggregate.  
**mirror** (STRING): The name of the mirror object that has aborted verification  
**duration** (STRING): The amount of time the verification required  
**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## **raid.mirror.verify.deferred**

### **Severity**

NOTICE

### **Description**

This message occurs when the mirror verification process is postponed due to inadequate resources. Mirror verification is a long running I/O operation that compares the blocks on both sides of mirror and reports any mismatches it finds. The operation will start when it get available resources.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s: verification deferred (%s)

### **Parameters**

**owner** (STRING): Owner of the affected aggregate.  
**mirror** (STRING): Name of the mirror object that could not be verified.  
**reason** (STRING): Reason code.

## **raid.mirror.verify.deferred.ok**

### **Severity**

NOTICE

### **Description**

This message occurs when a deferred mirror verify operation is now ready to proceed because of



availability of incore resources.

### Corrective Action

(None).

### Syslog Message

%s%s: verification previously deferred, now proceeding

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): Name of the mirror object that can now be verified.

## raid.mirror.verify.done

### Severity

NOTICE

### Description

This event is issued when verification has been completed on a specific mirror.

### Corrective Action

(None).

### Syslog Message

%s%s: verification completed in %s

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): The name of the mirror object that is completed verification

**duration** (STRING): The amount of time the verification required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.mirror.verify.mismatch

### Severity

NOTICE

### Description

This message occurs when mirror verification detected and corrected a mismatch.

### Corrective Action

(None).

### Syslog Message

%s: verify mismatch, disks %s and %s (vbn %llu, blockNum %llu)%s.

### Parameters

**grpName** (STRING): Name of the RAID group object that is being corrected.

**srcDisk** (STRING): Name of the disk with data used to correct the mismatch.

**dstDisk** (STRING): Name of the disk being corrected.

**volumeBno** (LONGINT): Volume block number.

**blockNum** (LONGINT): Disk block number.

**correcting** (STRING): " : correcting" if normal verify, "" if -n verify.

## raid.mirror.verify.resume

### Severity

NOTICE

### Description

This event is issued when a verify has been resumed on a mirror pair.

### Corrective Action

(None).

### Syslog Message

%s%s: resume mirror verification

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): The name of the mirror object that is resuming verification

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.mirror.verify.snapcrtfail

### Severity

ERROR

### Description

This message occurs when mirror verification fails to create a mirror verify Snapshot(tm) copy due to no space on the device for a Snapshot copy or the maximum number of Snapshot copies were reached. Mirror verification is a long running I/O operation that compares the blocks on both sides of mirror and reports any mismatches it finds.

### Corrective Action

Increase the size of the root volume or delete old Snapshot copies.

### Syslog Message

%s %s%s: could not create mirror verification Snapshot copy %s (%s)

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being verified.

**snapName** (STRING): Name of the Snapshot copy that was not created.

**error** (STRING): Error code returned by the failed operation.

## raid.mirror.verify.snapcrtok

### Severity

INFORMATIONAL

**Description**

This event is issued when verification has created a mirror verify snapshot.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s: created mirror verification snapshot %s

**Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being verified.

**snapName** (STRING): Name of the snapshot that has been created.

**raid.mirror.verify.snapdelfail****Severity**

NOTICE

**Description**

This message occurs when verification fails to delete a mirror verify Snapshot(tm) copy. The WAFL® Snapshot copy autodelete functionality will automatically delete the Snapshot copy.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s: could not delete mirror verification Snapshot copy %s (%s)

**Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate object that is being verified.

**snapName** (STRING): Name of the Snapshot copy that was not deleted.

**error** (STRING): Error code returned by the failed operation.

**raid.mirror.verify.snapdelok****Severity**

INFORMATIONAL

**Description**

This event is issued when verification has deleted a mirror verify snapshot.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s: deleted mirror verification snapshot %s

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate object that is being verified.  
**snapName** (STRING): Name of the snapshot that has been deleted.

## raid.mirror.verify.snaprenamefail

### Severity

NOTICE

### Description

This message occurs when verification fails to rename a mirror verify Snapshot(tm) copy. It is marked as invalid and is deleted.

### Corrective Action

(None).

### Syslog Message

%s %s%s: could not rename mirror verification Snapshot copy %s to %s (%s)

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate object that is being verified.  
**snapName** (STRING): Name of the Snapshot copy that was not renamed.  
**snapName2** (STRING): Attempted new name of the Snapshot copy.  
**error** (STRING): Error code returned by the failed operation.

## raid.mirror.verify.snaprenameok

### Severity

INFORMATIONAL

### Description

This event is issued when verification has renamed a mirror verify snapshot.

### Corrective Action

(None).

### Syslog Message

%s %s%s: renamed mirror verification snapshot %s to %s

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate object that is being verified.  
**snapName** (STRING): Old name of the snapshot that has been renamed.  
**snapName2** (STRING): New name of the snapshot that has been renamed.

## raid.mirror.verify.start

### Severity

NOTICE

### Description

This event is issued when a verify has been initiated on a mirror pair.

### Corrective Action

(None).

### Syslog Message

%s%s: start mirror verification

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): The name of the mirror object that is initiating verification

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.mirror.verify.suspend

### Severity

NOTICE

### Description

This event is issued when a verify has been suspended on a mirror pair.

### Corrective Action

(None).

### Syslog Message

%s%s: suspend mirror verification

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**mirror** (STRING): The name of the mirror object that is suspending verification

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.mirror.vote.badCksum

### Severity

INFORMATIONAL

### Description

This message occurs when the mirror vote blob contains an invalid checksum. The mirror vote blob holds information describing the active set of mirrored volumes.

### Corrective Action

(None).

## Syslog Message

RAID: mirror information has an inconsistent checksum.

## Parameters

(None).

## raid.mirror.vote.incorrectRecords

### Severity

INFORMATIONAL

### Description

This message occurs when the mirror vote blob has inconsistent contents. The mirror vote blob holds information describing the active set of mirrored volumes. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

RAID: mirror information has inconsistent contents.

## Parameters

**count** (INT): Number of records listed in the blob.

## raid.mirror.vote.invalidVersion

### Severity

INFORMATIONAL

### Description

This message occurs when the mirror vote blob contains an invalid version. The mirror vote blob holds information describing the active set of mirrored volumes.

### Corrective Action

(None).

## Syslog Message

RAID: mirror information has an unsupported version number (%d).

## Parameters

**version** (INT): Version information in the blob.

## raid.mirror.vote.invalidVol

### Severity

INFORMATIONAL

### Description

This message occurs when the mirror vote information describing a volume is internally inconsistent. The mirror vote information describes the active set of mirrored aggregate. The mirror vote information is

ignored.

### **Corrective Action**

(None).

### **Syslog Message**

RAID: mirror information for volume UUID %s is inconsistent.

### **Parameters**

**aggr\_id** (STRING): UUID of the affected volume.

## **raid.mirror.vote.noRecord**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a mirror vote is required on an aggregate but no vote is present in the mirror vote record. In this case, as both the plexes are available, we ignore mirror voting check.

### **Corrective Action**

(None).

### **Syslog Message**

RAID: mirror record missing for %s %s%s.

### **Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## **raid.mirror.vote.noRecord1Plex**

### **Severity**

ERROR

### **Description**

This message occurs when a mirror vote is required on an aggregate but no vote is present in the mirror vote record. In this case, only one plex is available. The aggregate will be kept offline because the existing plex might contain stale data.

### **Corrective Action**

1. In a storage environment with RAID SyncMirror®, the plex could be missing due to several reasons such as disaster at the site, shelf failure, disk failures, and so on. Try to bring the missing plex online after addressing some of the listed possible causes. 2. If the plex does not come online, bring the aggregate online using the 'storage aggr online' or 'storage plex online' command. This aggregate was kept offline because the existing plex might contain stale data.

### **Syslog Message**

WARNING: Only one plex in %s %s%s is available. %s might contain stale data.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**vol\_type2** (STRING): Volume type.

## raid.mirror.vote.outOfDate

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® detects an out-of-date plex. The plex is marked as out-of-date and a transaction subsequently occurs to update the internal RAID tree state.

### Corrective Action

(None).

### Syslog Message

%s %s%s has been detected as out-of-date and is being marked offline.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.

## raid.mirror.vote.sbFailed

### Severity

ERROR

### Description

This message occurs when the local node cannot transfer mirror vote records for aggregates being switched back as part of a switchback operation.

### Corrective Action

1. Verify that the DR partner node is up.
2. Run the 'network interface show' command to verify that the cluster network interfaces on the local and DR partner node are up. If not, correct any network issues that could be preventing it.

### Syslog Message

Could not communicate with the DR node over the intercluster network while attempting a switchback operation.

## Parameters

**dr\_host** (STRING): Disaster recovery (DR) host to which the local node failed to transfer the mirror vote records.

## raid.mirror.vote.versionZero



**Severity**

INFORMATIONAL

**Description**

This message occurs when the mirror vote blob contains version 0. Typically, this occurs when the mirror vote blob is empty.

**Corrective Action**

(None).

**Syslog Message**

RAID: mirror information is empty.

**Parameters**

(None).

**raid.mirror.vote.xferFailed****Severity**

ERROR

**Description**

This message occurs when the local node fails to transfer the mirror vote record of the aggregate during aggregate migration as part of giveback or aggregate relocation.

**Corrective Action**

1. Verify that the destination node is up.
2. Run the 'network interface show' command to verify that the cluster network interfaces on the local and partner node are up. If they are not up, address any network issues.

**Syslog Message**

Failed to communicate with the destination node over the cluster network while migrating the aggregate "%s" (UUID: %s) during giveback or aggregate relocation because %s.

**Parameters**

**aggregate** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): UUID of the aggregate.

**reason** (STRING): The reason the transfer failed.

**raid.monitor events****raid.monitor.maxVols****Severity**

ERROR

**Description**

This message occurs when the hourly monitor indicates that the current aggregate or volume number exceeded the maximum number of supported volumes on the host.

## Corrective Action

Check the number of volumes on the host. Destroy the newly created volume if the maximum for the total number of volumes on this host is exceeded.

## Syslog Message

WARNING: There are too many (%d) volumes on the host. The limit is %d in %s configuration.

## Parameters

**numVols** (INT): Current number of volumes on the host.  
**maxVols** (INT): Maximum number of volumes supported.  
**configType** (STRING): Current configuration.

# raid.multiperr events

## raid.multiperr.bad.block

### Severity

ERROR

### Description

This message occurs when RAID encounters more errors than the RAID level protection affords, and RAID cannot recover the block. The blocks with errors are marked as bad. When the file system reads this bad block, an error is returned. The file system identifies the bad file and block and recommends the necessary corrective action.

### Corrective Action

Check for the WAFL® message "waf.raid.incons.xxx" and follow the corrective actions outlined there.

### Syslog Message

Marking '%s%s', block number %llu, volume block number %llu, as a bad block.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): Physical block number containing the error.  
**volumeBno** (LONGINT): Volume block number.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multiperr.cksum.bno

## Severity

NOTICE

## Description

This message occurs when the system detects a block number mismatch during an error recovery operation. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).

## Syslog Message

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during an error recovery operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multierr.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during error recovery operations. The embedded checksum computed over the checksum entry does not match, indicating the corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s%s, block #%llu, during an error recovery operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multiperr.cksum.err

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum error during error recovery operations. The checksum computed does not match the stored checksum, indicating that the block is corrupted. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum error on %s%s, block #%llu during an error recovery operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multierr.cksum.rderr

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum block media error in an advanced\_zoned checksum's (AZCS) RAID group during an error recovery operation. In this case, the data cannot be verified. Data ONTAP® takes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during error recovery.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): First physical disk block number containing the error.  
**LblockNum** (LONGINT): Last physical disk block number containing the error.  
**shelf** (STRING): Shelf identifier for the disk shelf on which the disk is located.  
**bay** (STRING): Disk bay within the disk shelf on which the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multierr.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during an error recovery operation. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during an error recovery operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multiperr.lw.block.rewrite

### Severity

NOTICE

### Description

This message occurs on a RAID stripe with an inconsistent RAID write signature, when the system cannot detect the bad blocks or if the number of blocks with an error are more than the RAID protection level. The system restores the RAID write signature consistency by rewriting one or more disk blocks with the same data, but with the RAID write signature being corrected.

### Corrective Action

(None).

### Syslog Message

Rewriting %s%s, block #%llu with RAID write signature corrected.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): Physical Disk block number being rewritten.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multiperr.lw.block.rewrite.dirty

### Severity

INFORMATIONAL

## Description

This message occurs on a RAID stripe with an inconsistent RAID write signature that belongs to a dirty parity region, when the system cannot detect the bad blocks or if the number of blocks with an error are more than the RAID protection level. The system restores the RAID write signature consistency by rewriting one or more disk blocks with the same data, but with the RAID write signature being corrected.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): Physical Disk block number being rewritten.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.multiperr.lw.id.inconsist

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe while fixing multiple errors on a stripe. Data ONTAP® makes appropriate recovery actions. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

### Corrective Action

(None).

### Syslog Message

Inconsistent RAID write signature detected on RAID group %s%s, stripe #%llu, during RAID multiple error handling operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the raid group.

**stripe** (LONGINT): Stripe number.

## raid.multierr.lw.id.inconsist.dirty

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe while fixing multiple errors on a stripe that belongs to a dirty parity region.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

## raid.multierr.unverified.blk

### Severity

EMERGENCY

### Description

This message occurs when RAID encounters more errors than the RAID level protection allows and RAID cannot recover the block. Any block with a checksum error is marked as unverified.

### Corrective Action

Check for the WAFL® error message "wafl.raid.incons.xxx" and follow the corrective actions in that message, or contact NetApp technical support.

### Syslog Message

Marking '%s%s', block number [%llu - %llu], volume block number [%llu - %llu], as an unverified block.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): First physical block number containing the error.

**LblockNum** (LONGINT): Last physical block number containing the error.

**volumeBno** (LONGINT): First volume block number containing the error.

**LvolumeBno** (LONGINT): Last volume block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.



**site** (STRING): For a MetroCluster® configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.mv events

### raid.mv.defVol.online

#### Severity

INFORMATIONAL

#### Description

This message occurs when aggregates are brought online in a deferred way after mirror votes are retrieved automatically. Mirror votes are retrieved from the replicated database (RDB) in a deferred way when the RDB comes online.

#### Corrective Action

(None).

#### Syslog Message

RAID: Aggregate %s was brought online in a deferred way.

#### Parameters

**aggregate\_name** (STRING): Name of the aggregate that was brought online.

### raid.mv.defVol.online.fail

#### Severity

EMERGENCY

#### Description

This message occurs when the system fails to bring an aggregate online because mirror vote could not be retrieved from RDB. The reason could be a missing or incorrect mirror vote record or if the previous attempt to retrieve mirror votes failed. Mirror vote retrieval is retried automatically.

#### Corrective Action

The mirror vote retrieval is retried automatically. If that is not successful, aggregates that are kept offline due to the absence of mirror votes can be brought online by using the (privilege: diagnostic) ""storage aggregate online -forced"" command. Doing so, however, can cause data loss, and should only be attempted with help from NetApp Technical Support. Otherwise, wait until the other plex comes online for the aggregate to come online. For example, if one plex is missing due to disaster at that site, perform Disaster Recovery. Use the ""storage aggregate show"" command to see if plex is missing. If the condition persists, contact NetApp technical support.

#### Syslog Message

RAID: Could not bring aggregate %s online because %s.

#### Parameters

**aggregate\_name** (STRING): Name of the aggregate that could not be brought online.

**reason** (STRING): Reason the aggregate was not brought online.

## raid.mv.defVol.online.skip

### Severity

NOTICE

### Description

This message occurs when the system determines that the mirror vote should not be retrieved at this time, resulting in not bringing the aggregate online. The mirror vote retrieval process is retried automatically.

### Corrective Action

If the aggregate is a left-behind aggregate, use the "metrocluster switchback" command to return it to its original owner where it can be brought online. Otherwise, the system will automatically retry the mirror vote retrieval process.

### Syslog Message

RAID: Aggregate %s was not brought online because %s.

### Parameters

**aggregate\_name** (STRING): Name of the aggregate that was skipped.  
**reason** (STRING): Reason the aggregate was skipped.

## raid.mv.transfer.write

### Severity

NOTICE

### Description

This message occurs when mirror votes are received via HA message and are written to the local mailbox. Extended mirror vote write times can potentially prevent or delay the partner's ability to change ownership of a mirrored aggregate to this node.

### Corrective Action

(None).

### Syslog Message

The mailbox write for %d received records took %llu milliseconds.

### Parameters

**records** (INT): Number of new records written to the mailbox.  
**write\_time** (LONGINT): Time it took to write to the mailbox in milliseconds.

## raid.mv.unexpected.local

### Severity

INFORMATIONAL

### Description

This message occurs when we pull mirror votes from the replicated database (RDB), and there is a stale vote present in the mailbox. The vote from the RDB is treated as the correct vote, and the stale vote is discarded. The stale vote was probably due to not having been cleaned or deleted in the past.

## Corrective Action

(None).

## Syslog Message

RAID: Found a stale vote for aggregate %s in the mailbox when the mirror vote was pulled from the RDB.

## Parameters

**aggregate** (STRING): Name of the aggregate.

# raid.no events

## raid.no.parity.aggr

## Severity

EMERGENCY

## Description

This message occurs periodically to warn the user that the aggregate has no RAID protection and is not safe for storing data.

## Corrective Action

Use the "aggregate delete" command to delete the aggregate, if possible.

## Syslog Message

Aggregate "%s" does not use any RAID-level data protection; any data is unprotected.

## Parameters

**aggr** (STRING): Name of the aggregate.

# raid.ns events

## raid.ns.init.removeOwnership

## Severity

NOTICE

## Description

This message occurs when the system removes ownership of its disks during initialization of a non-shared HA system.

## Corrective Action

(None).

## Syslog Message

Removed ownership of %d disks.

## Parameters

**num\_disks\_reverted** (INT): Number of disks whose ownership was removed.

**num\_disks\_not\_reverted** (INT): Number of disks whose ownership was not removed.

**status** (STRING): Status of the disk ownership remove operation.

## raid.ns.init.skipUnpartition

### Severity

NOTICE

### Description

This message occurs when the system skips unpartitioning a disk during initialization of a non-shared HA system.

### Corrective Action

(None).

### Syslog Message

Skipped unpartitioning disk %s(UID: %s) due to %s.

### Parameters

**disk\_name** (STRING): Name of the disk.

**disk\_uid** (STRING): Unique identifier of the disk.

**reason** (STRING): Reason for skipping unpartitioning the disk.

## raid.ns.init.unpartition

### Severity

NOTICE

### Description

This message occurs when the system unpartitions one or more disks during initialization of a non-shared HA system.

### Corrective Action

(None).

### Syslog Message

Unpartitioned %d disk(s).

### Parameters

**num\_disks\_unpartitioned** (INT): Number of disks unpartitioned.

## raid.nvram events

### raid.nvram.bad.entry

### Severity

NOTICE

### Description

This message occurs when RAID NVRAM contains some bad entries. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains bad entries.

**Parameters**

(None).

**raid.nvram.clean.vol.entry****Severity**

NOTICE

**Description**

This message tracks the status of RAID nonvolatile RAM (NVRAM), which contained entries for a clean shutdown volume. No action is required.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains entries for a clean shutdown volume.

**Parameters**

(None).

**raid.nvram.miss.disk.entry****Severity**

NOTICE

**Description**

This message occurs when RAID NVRAM contains entries for nonexistent disks. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains entries for nonexistent disks.

**Parameters**

(None).

**raid.nvram.miss.group.entry****Severity**

NOTICE

**Description**

This message occurs when RAID NVRAM contains entries for nonexistent RAID groups. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains entries for nonexistent groups.

**Parameters**

(None).

**raid.nvram.miss.plex.entry****Severity**

NOTICE

**Description**

This message occurs when RAID NVRAM contains entries for nonexistent plexes. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains entries for nonexistent plexes.

**Parameters**

(None).

**raid.nvram.miss.vol.entry****Severity**

NOTICE

**Description**

This message occurs when RAID NVRAM contains entries for nonexistent volumes. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

RAID NVRAM contains entries for nonexistent volumes.

**Parameters**

(None).

## raid.nvram.stale.entry

### Severity

NOTICE

### Description

This message occurs when RAID NVRAM contains some stale entries. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

RAID NVRAM contains stale entries.

### Parameters

(None).

## raid.offline events

### raid.offline.busy

### Severity

NOTICE

### Description

This message occurs when an attempt to offline an aggregate fails due to active operations on the contained volumes.

### Corrective Action

(None).

### Syslog Message

Active operations in progress on aggregate %s: %s.

### Parameters

**aggrname** (STRING): Aggregate name.

**operation** (STRING): Operation.

## raid.olm events

### raid.olm.disk.ownership.change

### Severity

NOTICE

### Description

This message occurs when a disk which is under online label maintenance(OLM) is prepared for ownership change by administrator forcibly.

## Corrective Action

(None).

## Syslog Message

Ownership of %s, which is under online label maintenance (OLM) is being changed by administrator forcibly.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the disk shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.partition events

## raid.partition.disk

### Severity

NOTICE

### Description

This message indicates a disk was successfully partitioned.

### Corrective Action

(None).

### Syslog Message

Disk partition successful on %s, partitions created %u, partition sizes specified %u, partition spec summary%s.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**partition\_count** (INT): The number of partitions created.

**partition\_specs** (INT): The number of partitions created with the size specified.

**partition\_spec\_summary** (STRING): A summary of the partition sizes specified by partition specifications. The form is [index]=size [index]=size .. Where index is a partition index, and size capacity of the partition specified in disk blocks.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.



**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.partition.disk.fail

### Severity

NOTICE

### Description

This message occurs when an attempt to partition a disk fails. The error and error\_info fields indicate the reason why the attempt failed.

### Corrective Action

(None).

### Syslog Message

Disk partition failed on %s, error %s, additional error info (%s).

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**error** (STRING): String name for the error.

**error\_info** (STRING): Additional error information.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.plex events

### raid.plex.faultIsolation.reminder

### Severity

ERROR

### Description

This message occurs when disks in a plex do not have the same hardware-based pool.

### Corrective Action

a) Identify the problem aggregate from the system logs. b) Determine how this occurred: possible cause is a wiring problem. c) Based on the information you gather, determine how to correct the issue. For example, use the "storage disk replace" command to copy the disk belonging to the wrong pool to a disk belonging to the right pool. d) If you need assistance, contact NetApp technical support.

## Syslog Message

Plex %s has disks from mixed pools: %s

## Parameters

**plex** (STRING): Name of the plex.

**num\_disks** (STRING): Number of disks of the plex in other pools.

## raid.plex.vbn.error

### Severity

NOTICE

### Description

This message occurs when the system attempts to start up a plex and discovers that part of its Virtual Block Number(VBN) range is missing.

### Corrective Action

(None).

## Syslog Message

%s %s%s: Plex object %d is missing a VBN segment starting at %llu.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**plexId** (INT): Identifier of the plex.

**vbn** (LONGINT): Starting VBN that is missing.

## raid.pmdopt events

### raid.pmdOpt.misconfigured

### Severity

ERROR

### Description

The option `raid.panic.missing.disks` was set to `panic` and halt the appliance when a hardware error prevents access to multiple file system disks. However, the regular value of that option is not accessible during boot. Therefore, the option value is stored in the firmware environment variable `halt-missing-disks`, if the appliance has such capability. This event occurs hourly as long as the two values are not equal. Here are some examples of possible undesirable behavior caused by inconsistency between the two values: If `halt-missing-disks` is 0 and option `raid.panic.missing.disks` is 2, and if a failure causes the appliance to lose access to two or more file system disks, the appliance panics but does not halt. Multiple reconstructions might start after reboot, and that is what would have been prevented by setting the option. If `halt-missing-disks` is 2 and option `raid.panic.missing.disks` is 0, the appliance does not panic if it loses access to multiple file system disks. However, it might halt if multiple filesystem disks are missing during boot. If one value is 2 and the other is 3, the appliance panics and halts as expected only if a failure prevents access to three or more file system disks.

### Corrective Action

Try resetting the option `raid.panic.missing.disks` to the desired value.

### Syslog Message

The option `raid.panic.missing.disks` is set to %d, but during boot it will have value %d. Reset the option.

### Parameters

**value** (INT): Value of the option `raid.panic.missing.disks`.

**boot\_value** (INT): Value of the option `raid.panic.missing.disks` during boot that is stored in the firmware environment variable `halt-missing-disks`.

## raid.pmdOpt.nonpersistent

### Severity

NOTICE

### Description

This message occurs when an attempt is made to set the `"raid.panic.missing.disks"` option to `panic` and halt the system when a hardware error prevents access to multiple file system disks. However, the system cannot store the option persistently to be accessible during boot, and the option was not set to the new value. That happens on systems that do not have the capability to store environment variables in firmware. The `"raid.panic.missing.disks"` option cannot be used on such systems.

### Corrective Action

(None).

### Syslog Message

Setting the `"raid.panic.missing.disks"` option to %d failed because it could not be stored in the firmware environment variable.

### Parameters

**value** (INT): Value of the `"raid.panic.missing.disks"` option.

## raid.pmdOpt.set

### Severity

ERROR

### Description

The option `raid.panic.missing.disks` has been set to `panic` and halt the appliance when a hardware error prevents access to multiple file system disks. This option might be useful to prevent multiple unnecessary reconstructions if such errors occur, especially on NearStore® appliances. It is not recommended that you use this option on filer appliances. This option is always disabled during takeover. When using this option, you must have a monitoring system that alerts an administrator if the appliance halts. The Remote LAN Module (RLM) can provide such alerts if it is present in the appliance and properly configured. Otherwise, you need an external monitoring system.

### Corrective Action

Verify that the setting is correct. Verify that RLM or an external monitoring system can alert the administrator when the appliance halts. To disable this behavior, set option `raid.panic.missing.disks` to 0.

## Syslog Message

The option `raid.panic.missing.disks` is set to `panic` and halt the appliance when an error affects %d or more file system disks. RLM %s present.

## Parameters

**value** (INT): Value of the option `raid.panic.missing.disks`.

**rlm\_present\_string** (STRING): String to indicate whether RLM is present: "is" or "is not".

**rlm\_present** (INT): Indicates whether RLM is present in the system: 1 if RLM is present, 0 if it is not present.

## raid.pmdOpt.singleDisk

### Severity

ERROR

### Description

This message occurs when the option '`raid.panic.missing.disks`' is set to 1, which causes the system to panic and halt on any disk failure. You should set that option to 1 only when directed by NetApp technical support for handling a specific condition, and you should change that value from 1 as soon as the condition is rectified. This event is issued periodically as a reminder to change the option as soon as possible.

### Corrective Action

Change the option '`raid.panic.missing.disks`' to a more appropriate value (0, or 2 or higher) when conditions to keep it at 1 no longer apply.

## Syslog Message

The option '`raid.panic.missing.disks`' is set to 1, which causes the system to panic and halt on any disk failure. Change the option as soon as possible.

## Parameters

(None).

## raid.poolsort events

### raid.poolsort.disks\_per\_channel

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an invalid number of disks on a single channel.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**channel\_name** (STRING): Name of the channel containing too many disks.

**num\_disks** (INT): Number of disks connected to this channel.

**max\_disks** (INT): Maximum number of disks expected to be connected to this channel.

## **raid.poolsort.disks\_per\_shelf**

### **Severity**

ALERT

### **Description**

This message occurs when the system detects too many disks reporting the same disk shelf ID. This usually indicates that multiple shelves on one channel are using the same shelf ID.

### **Corrective Action**

Change the conflicting shelf IDs to unique values using the supported procedures for the affected shelves.

### **Syslog Message**

Shelf ID %d on channel %s is reporting too many disks.

### **Parameters**

**shelf\_id** (INT): Shelf ID being reported by the disks.  
**channel\_name** (STRING): Channel the disks are connected to.  
**num\_disks** (INT): Number of disks reporting this shelf ID.  
**max\_disks** (INT): Maximum number of disks expected to report this shelf ID.

## **raid.poolsort.num\_channels**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the system detects an invalid number of channels.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**num\_channels** (INT): Number of channels detected.  
**max\_channels** (INT): Maximum number of channels allowed on the system.

## **raid.poolsort.shelves\_per\_channel**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the system detects an invalid number of disk shelves on a single channel.

### **Corrective Action**

(None).

## Syslog Message

(None).

## Parameters

**channel\_name** (STRING): Channel containing too many shelves.

**num\_shelves** (INT): Number of shelves connected to this channel.

**max\_shelves** (INT): Maximum number of shelves expected to be connected to this channel.

# raid.preserve events

## raid.preserve.nvram

### Severity

INFORMATIONAL

### Description

During RAID replay of NVRAM, entries associated with a non-existent (offline or missing) aggregate were discovered. Since the volume is not available to replay those NVRAM entries, it will be saved in a system file on the root volume and will be replayed when the aggregate becomes available. This event is issued if the system is able to successfully save the NVRAM RAID entries associated with a non-existent aggregate.

### Corrective Action

(None).

### Syslog Message

RAID replay detected NVRAM entries for a non-existent aggregate having raidtreeID %s. %d checksum, %d stripe and %d pzero entries were successfully preserved.

### Parameters

**raidtreeID** (STRING): RAIDtree identifier of the aggregate.

**nchecksum\_saved** (INT): Identifies the number of checksum NVRAM entries preserved

**nstripe\_saved** (INT): Identifies the number of stripe NVRAM entries preserved

**nblkbuf\_saved** (INT): Identifies the number of pzero NVRAM entries preserved

## raid.preserve.nvram.issue

### Severity

INFORMATIONAL

### Description

This event occurs when system finds an issue during an nvfile save or nvfile replay operation. During RAID replay of NVRAM, entries associated with a nonexistent (offline or missing) aggregate are discovered.

Because the volume is unavailable to replay those NVRAM entries, it is saved in a system file on the root volume and is replayed when the aggregate becomes available.

### Corrective Action

(None).

### Syslog Message

%s %s

## Parameters

**reason** (STRING): Description of the issue during nvfile save or replay phase.

**identifier** (STRING): Description of raidtree identifier or filename

# raid.quiesce events

## raid.quiesce.diskProcessTime

### Severity

INFORMATIONAL

### Description

This message occurs during the ownership change of a disk to indicate the time taken by RAID to quiesce disk active I/O's.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**process** (STRING): Description of the process.

**disk** (STRING): Name of the disk.

**time** (LONGINT): Time taken by the specified process.

# raid.read events

## raid.read.cksum.blk.err

### Deprecated

Deprecated as of version 9.9.1 because this EMS does not appropriately describe the checksum block error.

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum block media error in an advanced zoned checksums(AZCS) RAID group during a read operation. Data ONTAP® takes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during RAID group I/O.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): First physical disk block number containing the error.

**LblockNum** (LONGINT): Last physical disk block number containing the error.

## raid.read.cksum.blk.notAvail

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum block media error in an advanced zoned checksums(AZCS) RAID group during a read operation. As a result checksum information is not available. Data ONTAP® takes appropriate recovery actions to reconstruct the block. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block is not available on %s%s for blocks [#%llu - #%llu] during RAID group I/O.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): First physical disk block number containing the error.

**LblockNum** (LONGINT): Last physical disk block number containing the error.

## raid.read.cksum.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch during a read operation. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during a read operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number,



vendor, model, RPM, and carrier serial number.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_dbn** (LONGINT): Expected physical disk block number.

**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected volume block number.

## raid.read.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during a read operation. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s%s, block #%llu, during a read operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical disk block number containing the error.

## raid.read.cksum.err

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum error during a read operation. The checksum computed does not match the stored checksum, indicating that the block is corrupted. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum error on %s%s, block #%llu.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical block number containing the error.

## raid.read.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during a read operation. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during a read operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical disk block number containing the error.

## raid.read.media.err

### Severity

NOTICE

### Description

This message occurs when rg i/o detects a single-disk medium error

### Corrective Action

(None).

### Syslog Message

Read error on %s%s, block #%llu

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.read.media.recommend.reassign.err

### Severity

NOTICE

### Description

This message occurs when rg i/o detects a single-disk recommend reassign error

### Corrective Action

(None).

### Syslog Message

Block recommended for reassignment on %s%s, block #%llu

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.readerr events

### raid.readerr.lw.data.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature, during a RAID error handling operation to log the on-disk RAID checksum information of the data block on the corresponding RAID stripe.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_name** (STRING): Name of the data disk.

**dbn** (LONGINT): Disk block number of the block being reported.  
**disk\_pos** (INT): Position of this data disk in the corresponding RAID group. The generation count of this data block on the parity bitmap resides at this position.  
**stripe\_id** (INTEX): Stripe ID field of the RAID write signature on the data block.  
**gen\_cnt** (INTEX): Generation count field of the RAID write signature on the data block.  
**comp\_cksum** (INTEX): Computed RAID checksum of the disk block.  
**sto\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**waf\_cxt0** (INTEX): Bytes 0-3 of WAFL® context information in the checksum.  
**waf\_cxt1** (INTEX): Bytes 4-7 of WAFL context information in the checksum.  
**waf\_cxt2** (INTEX): Bytes 8-11 of WAFL context information in the checksum.  
**waf\_cxt3** (INTEX): Bytes 12-15 of WAFL context information in the checksum.  
**waf\_cxt4** (INTEX): Bytes 16-19 of WAFL context information in the checksum.  
**waf\_cxt5** (INTEX): Bytes 20-23 of WAFL context information in the checksum.  
**waf\_cxt6** (INTEX): Bytes 24-27 of WAFL context information in the checksum.  
**waf\_cxt7** (INTEX): Bytes 28-31 of WAFL context information in the checksum.  
**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.lw.dparity.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature, during a RAID error handling operation to log the on-disk RAID checksum information of the diag parity block on the corresponding RAID stripe.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_name** (STRING): Name of the diag parity disk.  
**dbn** (LONGINT): Disk block number of the block being reported.  
**stripe\_id** (INTEX): Stripe ID field of the RAID write signature on the diag parity block.  
**gen\_cnt1** (INTEX): Generation count bitmap field of the RAID write signature on the diag parity block, for data blocks 0-15.  
**gen\_cnt2** (INTEX): Generation count bitmap field of the RAID write signature on the diag parity block, for data blocks 16-27.  
**comp\_cksum** (INTEX): Computed RAID checksum of the disk block.  
**waf\_cxt0** (INTEX): Bytes 0-3 of WAFL® context information in the checksum.  
**waf\_cxt1** (INTEX): Bytes 4-7 of WAFL context information in the checksum.  
**waf\_cxt2** (INTEX): Bytes 8-11 of WAFL context information in the checksum.  
**waf\_cxt3** (INTEX): Bytes 12-15 of WAFL context information in the checksum.  
**waf\_cxt4** (INTEX): Bytes 16-19 of WAFL context information in the checksum.  
**waf\_cxt5** (INTEX): Bytes 20-23 of WAFL context information in the checksum.  
**waf\_cxt6** (INTEX): Bytes 24-27 of WAFL context information in the checksum.  
**waf\_cxt7** (INTEX): Bytes 28-31 of WAFL context information in the checksum.  
**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.lw.id.inconsist

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe, during a RAID error handling operation. Data ONTAP® makes appropriate recovery actions. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

### Corrective Action

(None).

### Syslog Message

Inconsistent RAID write signature detected on RAID group %s%s, stripe #%llu, during RAID error handling operation

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

## raid.readerr.lw.id.inconsist.dirty

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe that belongs to a dirty parity region, during a RAID error handling operation.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

## raid.readerr.lw.parity.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature, during a RAID error handling operation to log the on-disk RAID checksum information of the row parity block on the corresponding RAID stripe.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_name** (STRING): Name of the row parity disk.

**dbn** (LONGINT): Disk block number of the block being reported.

**stripe\_id** (INTEHEX): Stripe ID field of the RAID write signature on the row parity block.

**gen\_cnt1** (INTEHEX): Generation count bitmap field of the RAID write signature on the row parity block, for data blocks 0-15.

**gen\_cnt2** (INTEHEX): Generation count bitmap field of the RAID write signature on the row parity block, for data blocks 16-27.

**comp\_cksum** (INTEHEX): Computed RAID checksum of the disk block.

**waf1\_cxt0** (INTEHEX): Bytes 0-3 of WAFL® context information in the checksum.

**waf1\_cxt1** (INTEHEX): Bytes 4-7 of WAFL context information in the checksum.

**waf1\_cxt2** (INTEHEX): Bytes 8-11 of WAFL context information in the checksum.

**waf1\_cxt3** (INTEHEX): Bytes 12-15 of WAFL context information in the checksum.

**waf1\_cxt4** (INTEHEX): Bytes 16-19 of WAFL context information in the checksum.

**waf1\_cxt5** (INTEHEX): Bytes 20-23 of WAFL context information in the checksum.

**waf1\_cxt6** (INTEHEX): Bytes 24-27 of WAFL context information in the checksum.

**waf1\_cxt7** (INTEHEX): Bytes 28-31 of WAFL context information in the checksum.

**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.lw.reconsBadBlk.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system encounters a RAID stripe with an inconsistent RAID write signature that might indicate a bad data block somewhere on the stripe. Based on this block's RAID signature, the system suspects this block as containing bad data and reconstructs the the block for a data integrity check. This message logs the RAID checksum information from the reconstructed block.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_name** (STRING): Name of the data disk.

**dbn** (LONGINT): Disk block number of the block being reported.

**comp\_cksum** (INTEHEX): Computed RAID checksum of the reconstructed block.

**waf1\_cxt0** (INTEHEX): Bytes 0-3 of WAFL® context information in the checksum.

**waf1\_cxt1** (INTEHEX): Bytes 4-7 of WAFL context information in the checksum.

**waf1\_cxt2** (INTEHEX): Bytes 8-11 of WAFL context information in the checksum.

**waf1\_cxt3** (INTEHEX): Bytes 12-15 of WAFL context information in the checksum.

**waf1\_cxt4** (INTHEX): Bytes 16-19 of WAFL context information in the checksum.  
**waf1\_cxt5** (INTHEX): Bytes 20-23 of WAFL context information in the checksum.  
**waf1\_cxt6** (INTHEX): Bytes 24-27 of WAFL context information in the checksum.  
**waf1\_cxt7** (INTHEX): Bytes 28-31 of WAFL context information in the checksum.  
**iteration** (INT): Parity consistency of this RAID stripe is restored by performing a data integrity check on one or two stripe blocks at a time in an iteration. This variable prints the current iteration.  
**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.lw.reconsGoodBlk.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system encounters a RAID stripe with an inconsistent RAID write signature that might indicate a bad data block somewhere on the stripe. Based on this block's RAID signature, the system identifies this block as containing good data and reconstructs the block to perform a data integrity check on a suspected bad data block in the stripe. This message logs the RAID checksum information from the reconstructed good block.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_name** (STRING): Name of the data disk.  
**dbn** (LONGINT): Disk block number of the block being reported.  
**comp\_cksum** (INTHEX): Computed RAID checksum of the reconstructed block.  
**waf1\_cxt0** (INTHEX): Bytes 0-3 of WAFL® context information in the checksum.  
**waf1\_cxt1** (INTHEX): Bytes 4-7 of WAFL context information in the checksum.  
**waf1\_cxt2** (INTHEX): Bytes 8-11 of WAFL context information in the checksum.  
**waf1\_cxt3** (INTHEX): Bytes 12-15 of WAFL context information in the checksum.  
**waf1\_cxt4** (INTHEX): Bytes 16-19 of WAFL context information in the checksum.  
**waf1\_cxt5** (INTHEX): Bytes 20-23 of WAFL context information in the checksum.  
**waf1\_cxt6** (INTHEX): Bytes 24-27 of WAFL context information in the checksum.  
**waf1\_cxt7** (INTHEX): Bytes 28-31 of WAFL context information in the checksum.  
**iteration** (INT): Parity consistency of this RAID stripe is restored by performing a data integrity check on one or two stripe blocks at a time in an iteration. This variable prints the current iteration.  
**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.lw.tparity.ckinfo

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature, during a RAID error handling operation to log the on-disk RAID checksum information of the tparity block on the corresponding RAID stripe.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_name** (STRING): Name of the tparity disk.

**dbn** (LONGINT): Disk block number of the block being reported.

**stripe\_id** (INTHEX): Stripe ID field of the RAID write signature on the tparity block.

**gen\_cnt1** (INTHEX): Generation count bitmap field of the RAID write signature on the tparity block, for data blocks 0-15.

**gen\_cnt2** (INTHEX): Generation count bitmap field of the RAID write signature on the tparity block, for data blocks 16-27.

**comp\_cksum** (INTHEX): Computed RAID checksum of the disk block.

**waf1\_cxt0** (INTHEX): Bytes 0-3 of WAFL® context information in the checksum.

**waf1\_cxt1** (INTHEX): Bytes 4-7 of WAFL context information in the checksum.

**waf1\_cxt2** (INTHEX): Bytes 8-11 of WAFL context information in the checksum.

**waf1\_cxt3** (INTHEX): Bytes 12-15 of WAFL context information in the checksum.

**waf1\_cxt4** (INTHEX): Bytes 16-19 of WAFL context information in the checksum.

**waf1\_cxt5** (INTHEX): Bytes 20-23 of WAFL context information in the checksum.

**waf1\_cxt6** (INTHEX): Bytes 24-27 of WAFL context information in the checksum.

**waf1\_cxt7** (INTHEX): Bytes 28-31 of WAFL context information in the checksum.

**disk\_serialno** (STRING): Serial number of the disk drive.

## raid.readerr.recovered.file

### Severity

NOTICE

### Description

This event is issued when a lost write block has been recovered from parity and other remaining data disks.

### Corrective Action

(None).

### Syslog Message

Recovered the bad disk block %llu on %s%s, %s inode number %d, snapid %d, treeid %d, file block %llu level %d.

### Parameters

**blockNum** (LONGINT): The disk block number

**owner** (STRING): String indicating the owner of the affected volume.

**disk\_info** (STRING): Formatted information of the disk

**ino\_type** (STRING): The inode type

**fileid** (INT): The file identifier

**snapid** (INT): The snapshot identifier

**treeid** (INT): The snapshot identifier

**fileBlockNum** (LONGINT): The block number in the file

**level** (INT): The level



# raid.recons events

## raid.recons.lw.id.inconsist

### Severity

NOTICE

### Description

This message occurs during a RAID reconstruction operation, when the system detects an inconsistent RAID write signature on a RAID stripe. Data ONTAP® makes appropriate recovery actions. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

### Corrective Action

(None).

### Syslog Message

Inconsistent RAID write signature detected on RAID group %s%s, stripe #%llu, during reconstruction.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the raid group.

**stripe** (LONGINT): Stripe number.

## raid.recons.lw.id.inconsist.dirty

### Severity

INFORMATIONAL

### Description

This message occurs during a RAID reconstruction operation, when the system detects an inconsistent RAID write signature on a RAID stripe that belongs to a dirty parity region.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

## raid.recons.MP

### Severity

NOTICE

## Description

This message occurs when a user enables or disables multiprocessing for reconstruction. It also is emitted when reconstruction actually starts and multiprocessing is enabled.

## Corrective Action

(None).

## Syslog Message

RAID reconstruct multiprocessing is %s (raid.reconstruct.threads: %d rows, RAID I/O threads: %d).

## Parameters

**status** (STRING): Whether multiprocessing is enabled or disabled

**rows** (INT): Number of reconstruct rows ("raid.reconstruct.threads") in use per RAID I/O thread.

**raidio\_nthreads** (INT): Number of RAID I/O threads in use.

# raid.reparity events

## raid.reparity.lw.id.inconsist

### Severity

INFORMATIONAL

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe, during a RAID reparity operation.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

# raid.replay events

## raid.replay.nvfile.done

### Severity

INFORMATIONAL

### Description

This event is issued when an aggregate becomes available and system successfully replayed the preserved NVRAM entries associated with that aggregate. The NVRAM entries replayed are saved in a system file on the root volume.

## Corrective Action

(None).

## Syslog Message

Successfully replayed %d checksum, %d stripe %d pzero block entries from preserved NVRAM file for %s %s%s having raidtreeID %s.

## Parameters

**nchecksum\_replayed** (INT): Identifies the number of checksum NVRAM entries replayed

**nstripe\_replayed** (INT): Identifies the number of stripe NVRAM entries replayed

**nblkbuf\_replayed** (INT): Identifies the number of pzero entries replayed

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**raidtreeID** (STRING): RAIDtree identifier of the aggregate.

## raid.replay.nvfile.start

### Severity

INFORMATIONAL

### Description

This event is issued while starting the raid nvfile replay for an aggregate. When a aggregate becomes available and system determines that there are relevant nvram entries saved in a system file for such aggregate, those saved nvram entries must be replayed.

### Corrective Action

(None).

### Syslog Message

Starting raid nvfile replay for %s %s%s having raidtreeID %s stored on %s host root volume.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**raidtreeID** (STRING): Raidtree identifier of the aggregate.

**host** (STRING): The preserved nvfile being replayed for this aggregate was found on the host root volume.

## raid.replay.partner.nvram

### Severity

NOTICE

### Description

This message occurs when RAID replays partner nonvolatile RAM (NVRAM).

### Corrective Action

(None).

## Syslog Message

Replaying partner NVRAM.

## Parameters

(None).

# raid.rg events

## raid.rg.bigio.fatal

### Severity

ALERT

### Description

This message occurs when a RAID group background long running I/O encounters an error that is fatal to the RAID group and the aggregate (cannot be fixed even by a mirror read from the other plex, if it exists).

### Corrective Action

Contact technical support to bring this aggregate online.

### Syslog Message

%s: long-running raid I/O operation experienced fatal multi-disk error.

### Parameters

**rg** (STRING): Name of the RAID group object that failed.

## raid.rg.context.clear

### Severity

NOTICE

### Description

This message occurs when the context flag is cleared from the RAID group, and therefore from the parent aggregate, as part of context parity scrub testing. This event is only logged during internal testing; it is not seen on customer systems.

### Corrective Action

(None).

### Syslog Message

Scrub testing: %s%s: group parity context is now reset.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group whose context flag is cleared.

## raid.rg.degraded

### Severity

NOTICE

## Description

This message occurs when a RAID group becomes degraded, double degraded, or triple degraded.

## Corrective Action

Reconstruction will occur unless there is no matching spare disk or the number of failed disks is greater than the parity reconstruction level of the RAID group (1 for raid4, 2 for raid\_dp and 3 for raid\_tec). If reconstruction does not occur, then follow the following steps: 1. Identify the failed disk or disks from the system logs and determine the reason for the failure. 2. Replace the faulty disks.

## Syslog Message

%s: Raid group %s is degraded

## Parameters

**owner** (STRING): Owner of the RAID group.

**name** (STRING): Relative path name of the RAID group.

**aggr\_UUID** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.die\_recons.aborted

### Severity

NOTICE

### Description

This message occurs when die reconstruction has been aborted on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: die reconstruction aborted at disk %s block %llu after %s

### Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is aborted die reconstruction.

**disk** (STRING): Name of the disk object for which die reconstruction aborted.

**blockNum** (LONGINT): The physical block number die reconstruction stopped

**duration** (STRING): The amount of time the die reconstruction required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.die\_recons.cant.start

### Severity

NOTICE

### Description

This message occurs when the raid group has a disk with failed die but the die reconstruction cannot be started.

### Corrective Action

(None).

## Syslog Message

%s%s: unable to start die reconstruction for %s: %s.

## Parameters

**owner** (STRING): Owner of the raid group  
**rg** (STRING): Name of the raid group  
**disk** (STRING): Name of the disk which has a failed die  
**reason** (STRING): Reason the die reconstruction cannot start

## raid.rg.die\_recons.deferred

### Severity

NOTICE

### Description

This message occurs when die reconstruction has been deferred on a specific raid group due to inadequate incore resources.

### Corrective Action

(None).

## Syslog Message

%s%s: die reconstruction deferred for %s (%s)

## Parameters

**owner** (STRING): Owner of the raid group  
**rg** (STRING): Name of the raid group object that is unable to reconstruct die  
**disk** (STRING): Name of the disk object for which die reconstruction aborted.  
**reason** (STRING): A text reason code

## raid.rg.die\_recons.done

### Severity

NOTICE

### Description

This message occurs when die reconstruction has been completed on a specific raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: die reconstruction completed for %s in %s.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.  
**owner** (STRING): Owner of the raid group.  
**disk** (STRING): Name of the disk object for which die reconstruction completed.  
**duration** (STRING): The amount of time the die reconstruction required.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.  
**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.die\_recons.start

### Severity

NOTICE

### Description

This message occurs when die reconstruction has been initiated on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: starting die reconstruction on disk %s

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.  
**owner** (STRING): Owner of the raid group.  
**disk** (STRING): Name of the disk object where die reconstruction happens  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.disk.reconstruction.failed

### Severity

NOTICE

### Description

This event is issued when a disk has failed reconstruction.

### Corrective Action

(None).

## Syslog Message

%s%s: reconstruction failed for a disk in the raidgroup

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object containing the reconstructing disk

## raid.rg.diskcopy.aborted

### Severity

NOTICE

### Description

This message occurs when RAID disk copy is aborted. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

%s%s: disk copy from %s to %s aborted at disk block %llu after %s. Reason: %s.

## Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): RAID group name.

**source** (STRING): Name of the source disk.

**target** (STRING): Name of the target disk.

**blockNum** (LONGINT): Physical block number where the disk copy stopped.

**duration** (STRING): Amount of time the disk copy lasted.

**reason** (STRING): Reason for aborting the disk copy operation

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.diskcopy.cant.start

### Severity

NOTICE

### Description

This message occurs when RAID disk copy cannot be started. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

%s%s: unable to start disk copy for %s: %s.

## Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

**disk** (STRING): Name of the disk to be copied.



**reason** (STRING): Reason the disk copy cannot start.

## **raid.rg.diskcopy.deferred**

### **Severity**

NOTICE

### **Description**

This message occurs when a RAID disk copy is deferred due to inadequate incore resources. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s: disk copy for %s deferred (%s).

### **Parameters**

**owner** (STRING): Owner of the RAID group.  
**rg** (STRING): Name of the RAID group.  
**disk** (STRING): Name of the disk to be copied.  
**reason** (STRING): Reason code.

## **raid.rg.diskcopy.deferred.ok**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a previously deferred RAID disk copy is now ready to proceed.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s: disk copy for %s previously deferred, now proceeding.

### **Parameters**

**owner** (STRING): Owner of the RAID group.  
**rg** (STRING): Name of the RAID group.  
**disk** (STRING): Name of the disk to be copied.

## **raid.rg.diskcopy.done**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when RAID disk copy is completed.

## Corrective Action

(None).

## Syslog Message

%s%s: disk copy from %s (S/N [%s]) to %s (S/N [%s]) was completed in %s.

## Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): RAID group name.

**source** (STRING): Name of the source disk.

**source\_serialno** (STRING): Serial number of the source disk.

**target** (STRING): Name of the target disk.

**target\_serialno** (STRING): Serial number of the target disk.

**duration** (STRING): Amount of time it took to complete the disk copy operation.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.diskcopy.failed

### Severity

NOTICE

### Description

This message occurs when RAID disk copy to the intended target disk cannot be initiated. This type of disk copy is used after new firmware download, to return the disk to its place in a RAID group. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s%s: unable to start disk copy from %s to %s: %s.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

**source** (STRING): Name of the disk to be copied.

**target** (STRING): Name of the target disk.

**reason** (STRING): Reason the disk copy cannot start.

## raid.rg.diskcopy.mismatch

### Severity

ERROR

### Description

This message occurs when there is no disk for Rapid RAID Recovery that matches all parameters of an ideal replacement in the RAID group. Data ONTAP® selected a disk with a different interface speed or RPM, or from the opposite spare pool in a SyncMirror® configuration.

### Corrective Action

Wait for copying to finish. Install an appropriate spare disk. Fail the mismatched disk by using the 'disk fail' command. It will be copied to the new spare disk. After copying is done, use the 'disk unfail' command to

unfail the mismatched disk and convert it to a spare.

### Syslog Message

%s%s: matching disk not available to copy disk %s; using disk %s %s.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

**source** (STRING): Name of the source disk.

**target** (STRING): Name of the target disk.

**detail** (STRING): Description of disk mismatch; one or more of: from the opposite spare pool, with different RPM.

## raid.rg.diskcopy.recom.fail

### Severity

ERROR

### Description

This message occurs when an unhealthy disk drive being copied to a replacement disk cannot continue the copy, and must be recommended for failure before the copy is completed. The system will reconstruct the data from parity data.

### Corrective Action

(None).

### Syslog Message

%s recommend failed after exceeding disk copy error threshold.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**error\_pct** (INT): Number of bad blocks on the disk, described as percentage of the total size of the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.diskcopy.resume

### Severity

INFORMATIONAL

### Description

This message occurs when RAID disk copy is resumed.

## Corrective Action

(None).

## Syslog Message

%s%s: resuming disk copy from %s to %s (disk block %llu, %d%% complete)

## Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): RAID group name.

**source** (STRING): Name of the source disk.

**target** (STRING): Name of the target disk.

**blockNum** (LONGINT): Physical block number from which disk copy is continuing.

**percent** (INT): Percentage of disk copy already complete.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.diskcopy.start

### Severity

NOTICE

### Description

This message occurs when RAID disk copy is initiated.

### Corrective Action

(None).

### Syslog Message

%s%s: starting disk copy from %s (S/N [%s]) to %s (S/N [%s]). Reason: %s.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): RAID group name.

**source** (STRING): Name of the source disk.

**source\_serialno** (STRING): Serial number of the source disk.

**target** (STRING): Name of the target disk.

**target\_serialno** (STRING): Serial number of the target disk.

**reason** (STRING): Reason for starting disk copy.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.drecons

### Severity

NOTICE

### Description

This message occurs when double reconstruction is initiated after the failure of two disks in RAID-DP or RAID-TEC raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: starting double reconstruction, using disks %s and %s

## Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is initiating double reconstruction.

**disk** (STRING): Name of the disk object that is the first reconstruction target.

**disk1** (STRING): Name of the disk object that is the second reconstruction target.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.intl.diskcopy.done

### Severity

NOTICE

### Description

This message occurs when intelligent disk copy has been completed.

### Corrective Action

(None).

## Syslog Message

%s%s: intelligent disk copy completed for %s. Used DBN %llu copy done DBN %llu disk size %llu

## Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is completed intelligent disk copy.

**disk** (STRING): Name of the disk object for which intelligent reconstruction completed.

**used\_dbn** (LONGINT): Maximum used Disk block number.

**copy\_done\_dbn** (LONGINT): Disk block number when intelligent disk copy was done.

**disk\_size** (LONGINT): Filesystem size of the disk.

## raid.rg.intl.recons.done

### Severity

NOTICE

### Description

This message occurs when intelligent reconstruction has been completed on a specific raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: intelligent reconstruction completed for %s. Used DBN %llu reconstruction done DBN %llu disk size %llu

## Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is completed intelligent reconstruction.

**disk** (STRING): Name of the disk object for which intelligent reconstruction completed.

**used\_dbn** (LONGINT): Maximum used Disk block number.  
**recons\_done\_dbn** (LONGINT): Disk block number when intelligent reconstruction was done.  
**disk\_size** (LONGINT): Filesystem size of the disk.

## **raid.rg.lw.prot.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the lost-write protection for RAID I/O is disabled on the RAID group.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s: group lost-write protection for RAID I/O is now disabled (%s)

### **Parameters**

**owner** (STRING): Owner of the RAID group.  
**rg** (STRING): Name of the RAID group on which lost-write protection is disabled.  
**reason** (STRING): A character string indicating the reason for disabling the lost-write protection for RAID I/O.

## **raid.rg.media\_scrub.cksum\_blk.err.bypass**

### **Severity**

NOTICE

### **Description**

For a zoned checksum aggregate this event is issued when background media scrub bypasses a media error on a checksum block.

### **Corrective Action**

(None).

### **Syslog Message**

Media error on checksum block bypassed for %s%s, block #%llu

### **Parameters**

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): The physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.media\_scrub.done

### Severity

NOTICE

### Description

This event is issued when media scrub has completed on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: media scrub completed in %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that has completed media scrub

**duration** (STRING): The amount of time the media scrub required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.media\_scrub.media.err

### Severity

NOTICE

### Description

This event is issued when a media\_scrub i/o detects a medium error

### Corrective Action

(None).

### Syslog Message

Read error on %s%s, block #%llu during media\_scrub

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): The name of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.media\_scrub.media.errs

### Severity

NOTICE

### Description

This message occurs when background media scrubbing detects a multi-disk media error on a stripe. Data ONTAP® will take appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Multiple checksum block read error on %s%s for blocks [#%llu - #%llu] during media\_scrub.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group object that contains the multi-disk media error.

**blockNum** (LONGINT): First physical disk block number containing the error.

**LblockNum** (LONGINT): Last physical disk block number containing the error.

## raid.rg.media\_scrub.media.recommend.reassign.err

### Severity

NOTICE

### Description

This event is issued when a background media scrub i/o detects a recommend reassign error

### Corrective Action

(None).

### Syslog Message

Block recommended for reassignment on %s%s, block #%llu during media\_scrub

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.



## raid.rg.media\_scrub.resume

### Severity

NOTICE

### Description

This event is issued when media\_scrub resumes on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: resuming media scrub at stripe %llu (%d%% complete)

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is resuming media\_scrub

**stripe** (LONGINT): resuming at stripe.

**percentage** (INT): (None).

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.media\_scrub.rewrite.chunk

### Severity

INFORMATIONAL

### Description

This message occurs when background media scrub rewrites a chunk of blocks for ATA drives.

### Corrective Action

(None).

### Syslog Message

media\_scrub rewriting chunk on %s%s, start block #%llu, count #%u.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical block number at which the chunk starts.

**count** (INT): Number of blocks being rewritten.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.media\_scrub.start

### Severity

NOTICE

### Description

This event is issued when media\_scrub has been initiated on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: start media scrub

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is initiating media\_scrub

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.media\_scrub.stopped

### Severity

NOTICE

### Description

This event is issued when media scrub is stopped on specified raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: media scrub stopped at stripe %llu after %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object for which media scrub was stopped

**stripe** (LONGINT): stopped at stripe.

**duration** (STRING): The amount of time the media scrub has taken.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.media\_scrub.summary.media

### Severity

NOTICE

### Description

This event is generated as a summary report for errors found during a media scrub.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**errors** (INT): The total number of media errors found

**rg** (STRING): The name of the raid group

**current** (STRING): The current status

## raid.rg.media\_scrub.suspended

### Severity

NOTICE

### Description

This event is issued when media scrub is suspended on specified raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: media scrub suspended at stripe %llu after %s

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object for which media scrub is suspended

**stripe** (LONGINT): suspended at stripe.

**duration** (STRING): The amount of time the media scrub has taken.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.mirrorreaderr.fatal

### Severity

NOTICE

### Description

This message occurs when the RAID mirror read error handling encounters an I/O error on the RAID group in the source plex. This is ultimately fatal error to the aggregate. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

## Syslog Message

%s: RAID mirror read error handling experienced fatal multi-disk error; failing RAID group.

## Parameters

**rg** (STRING): Name of the RAIDgroup that failed.

## raid.rg.mirrorreaderr.nonfatal

**Severity**

NOTICE

**Description**

This message occurs when the RAID mirror read error handling encounters an I/O error on a RAID group during a scrub, or lost write recovery operation. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

%s: RAID mirror read error handling experienced multi-disk error; continuing.

**Parameters**

**rg** (STRING): Name of the RAID group that has a double error.

**raid.rg.normal****Severity**

NOTICE

**Description**

This message occurs when a RAID group returns to normal after it was in one of the degraded states, after a resynchronization, or after it was in failed state.

**Corrective Action**

(None).

**Syslog Message**

%s: raid group %s is in normal state.

**Parameters**

**owner** (STRING): Owner of the raid group.

**name** (STRING): Relative path name of the raid group.

**aggr\_UUID** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**raid.rg.readerr.bad.file.block****Severity**

EMERGENCY

**Description**

This message occurs when a reconstructed file block is found out to be marked bad because it could not be recovered correctly earlier.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Reading bad block on %s %s%s, %s inode number %d, snapid %d, file block %llu, level %d.

## Parameters

**vol\_type** (STRING): String indicating whether the file is contained in a flexible volume or aggregate.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the volume or aggregate  
**ino\_type** (STRING): The inode type  
**fileid** (INT): File identifier  
**snapid** (INT): Snapshot identifier  
**fileBlockNum** (LONGINT): Block number in the file  
**level** (INT): Inode indirect pointer level

## raid.rg.readerr.block.fail

### Severity

ERROR

### Description

This event is issued when a missing block could not be properly recovered because a read error that was detected on another disk in the group was bypassed.

### Corrective Action

Ensure wafiron successfully completes. Wafiron will preserve the filesystem consistency.

### Syslog Message

Block #%llu cannot be properly recovered on %s%s

## Parameters

**blockNum** (LONGINT): The physical block number that could not be recovered  
**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk to which the block belongs  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.recalc

### Severity

NOTICE

### Description

This event is issued when a previously detected bad block is recalculated.

### Corrective Action

(None).

## Syslog Message

Recalculated bad block on %s%s, block #%llu

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): The name of the disk  
**blockNum** (LONGINT): The disk block number  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.recommend.failure

### Severity

ERROR

### Description

This message occurs when a detected bad block is repeatedly recalculated and written back to disk.

### Corrective Action

Wait for the event message `raid.disk.unload.done` or `raid.carrier.remove` to be issued and then replace the carrier containing the failed disk. If the disk is contained in a single-disk carrier, `raid.disk.unload.done` will be issued immediately. Otherwise, wait for `raid.carrier.remove` which indicates the carrier has been fully prepared for removal.

## Syslog Message

Disk failure due to repeated read errors on %s%s, block #%llu

## Parameters

**owner** (STRING): String indicating the owner of the affected volume.  
**disk\_info** (STRING): The name of the disk.  
**blockNum** (LONGINT): The physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.cksum.computed

### Severity

NOTICE

### Description

This message occurs when a previously detected bad block is repaired and the checksum of that recalculated block matches the computed checksum.

### Corrective Action

(None).

### Syslog Message

Recalc %s%s, block #%llu matches computed checksum

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk  
**blockNum** (LONGINT): The disk block number  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.cksum.error

### Severity

NOTICE

### Description

This message occurs when a previously detected bad block is repaired and the checksum of that recalculated block does not match the stored and computed checksums.

### Corrective Action

(None).

### Syslog Message

Data recalculated from parity on %s%s; block #%llu does not match the stored and computed checksums.

### Parameters

**owner** (STRING): String indicating the owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier ID.  
**blockNum** (LONGINT): Disk block number.  
**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.cksum.stored

### Severity

NOTICE

### Description

This message occurs when a previously detected bad block is repaired and the checksum of that recalculated block matches the stored checksum.

### Corrective Action

(None).

### Syslog Message

Recalc %s%s, block #%llu matches stored checksum

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk  
**blockNum** (LONGINT): The disk block number  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.data

### Severity

NOTICE

### Description

This message occurs when fixing a single-disk media or checksum error on a data disk.

### Corrective Action

(None).



## Syslog Message

Fixing bad data on %s%s, disk block (DBN) %llu, aggregate block (VBN) %llu.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number containing the block being rewritten.  
**blockNum** (LONGINT): Disk block number.  
**vbn** (LONGINT): Volume block number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.data.dirty

### Severity

NOTICE

### Description

This message occurs when fixing a single-disk media or checksum error on a data block that belongs to a dirty parity region.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number containing the block being rewritten.  
**blockNum** (LONGINT): Disk block number.  
**vbn** (LONGINT): Volume block number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.filebno

### Severity

NOTICE

### Description

This event is issued when a previously detected bad block is repaired.

### Corrective Action

(None).

### Syslog Message

%s %s%s, %s inode number %d, snapid %d, treeid %d, uid %u, gid %u, file block %llu level %d

### Parameters

**vol\_type** (STRING): Whether this error occurred in an aggregate, or a flexible volume.

**owner** (STRING): Owner of the affected volume or aggregate.

**vol** (STRING): The name of the volume or aggregate

**ino\_type** (STRING): The inode type

**fileid** (INT): The file identifier

**snapid** (INT): The snapshot identifier

**treeid** (INT): The qtree identifier

**uid** (INT): The user id that owns the file

**gid** (INT): The group id that owns the file

**bno** (LONGINT): The block number in the file

**level** (INT): The level of the block within the file

## raid.rg.readerr.repair.parity

### Severity

NOTICE

### Description

This message occurs when fixing a single-disk media or checksum error on a parity disk.

### Corrective Action

(None).

### Syslog Message

Fixing bad parity on %s%s, block #%llu

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk containing the block being rewritten.

**blockNum** (LONGINT): The disk block number

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.repair.parity.dirty

### Severity

NOTICE

### Description

This message occurs when fixing a single-disk media or checksum error on a parity block that belongs to a dirty parity region.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk containing the block being rewritten.

**blockNum** (LONGINT): The disk block number

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.readerr.wc.blkErr

### Severity

EMERGENCY

### Description

This event occurs when a checksum error due to WAFL context mismatch is detected on a reconstructed block. This error indicates data loss and are caused by misbehaving disks.

### Corrective Action

Start wafliion on the indicated volume or aggregate.

### Syslog Message

Checksum error due to wafli context mismatch on %s %s%s, %s inode number %d, snapid %d, file block %llu, level %d, RBN %llu. Checksum context has FBN %llu, buftreeid %u, CP count %u, encrypted flag %d, key index %d, PVBN %lld.

## Parameters

**vol\_type** (STRING): String indicating whether the file is contained in a flexible volume or aggregate.  
**owner** (STRING): Owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate  
**ino\_type** (STRING): The inode type  
**fileid** (INT): File identifier  
**snapid** (INT): Snapshot identifier  
**fileBlockNum** (LONGINT): Block number in the file  
**level** (INT): Inode indirect pointer level  
**blockNum** (LONGINT): RAID Block Number (RBN) of the block  
**stored\_fbn** (LONGINT): Stored file block number in the checksum entry.  
**stored\_bufreeid** (INT): Stored buftree identifier in the checksum entry.  
**stored\_cpcount** (INT): Stored CP count in the checksum entry.  
**stored\_encrypted\_flag** (INT): Stored encrypted flag in the checksum entry.  
**stored\_key\_index** (INT): Stored encryption key index in the checksum entry.  
**pvpn** (LONGINT): Physical Volume Block Number (PVPN) of the block

## raid.rg.recon.bad.cksum.embd

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during reconstruction while the aggregate is in media error bypass mode. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. The error is ignored.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s%s, block #%llu during reconstruction while in media error bypass mode.

## Parameters

**owner** (STRING): Owner of the disk drive.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recon.bad.cksum.zero

## Severity

NOTICE

## Description

This message occurs when the system detects an empty checksum entry during reconstruction while the aggregate is in media error bypass mode. The checksum entry is zeroed, but the corresponding block is not zeroed. The error is ignored.

## Corrective Action

(None).

## Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during reconstruction while in media error bypass mode.

## Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): Physical disk block number containing the error.

**shelf** (STRING): Disk shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.aborted

## Severity

NOTICE

## Description

This message occurs when reconstruction has been aborted on a specific raid group.

## Corrective Action

(None).

## Syslog Message

%s%s: reconstruction aborted at disk block %llu after %s

## Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is aborted reconstruction.

**blockNum** (LONGINT): The physical block number reconstruction stopped

**duration** (STRING): The amount of time the reconstruction required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.recons.bad.cksum

### Severity

NOTICE

### Description

This event is issued when reconstruction detects a bad checksum on the stripe while the aggregate is in media error bypass mode.

### Corrective Action

(None).

### Syslog Message

Checksum mismatch on %s%s, block #%llu during reconstruction; assuming data is correct and fixing checksum

### Parameters

**owner** (STRING): Owner of the aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.bad.cksum.blk

### Severity

NOTICE

### Description

This message occurs when reconstruction corrects a checksum block media error in an advanced\_zoned checksums (azcs) RAID group on a stripe while in media error bypass mode. Data ONTAP® takes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during reconstruction; assuming data is correct and fixing checksum.

### Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information about the disk object that contains the error.

**blockNum** (LONGINT): First physical disk block number containing the error.  
**LblockNum** (LONGINT): Last physical disk block number containing the error.  
**shelf** (STRING): Disk Shelf identifier for the shelf on which the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf on which the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.bad.cksum.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch during reconstruction while the aggregate is in media error bypass mode. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. The error is ignored.

### Corrective Action

(None).

### Syslog Message

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during reconstruction while in media error bypass mode.

### Parameters

**owner** (STRING): Owner of the disk drive.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Disk Shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.block.fail

### Severity

ERROR

### Description

A block on the reconstructing disk can not be reconstructed because of a media error on another disk in the group. The block is zeroed and written to disk. If the block is allocated, data in that block is corrupted.

### Corrective Action

Ensure wafiron successfully completes after this reconstruction is finished. Wafiron will preserve the filesystem consistency.

### Syslog Message

Block #%llu can not be properly reconstructed on %s%s - the block will be zeroed

### Parameters

**blockNum** (LONGINT): The physical block number that could not be reconstructed

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk being reconstructed

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.cantStart

### Severity

ERROR

### Description

This event is issued when a RAID group is missing a disk but the reconstruction cannot be started.

### Corrective Action

Usually, an appropriate spare disk should be installed. Sometimes, the quoted reason and presence of the appropriate spare disks can indicate a different problem.

### Syslog Message

The reconstruction cannot start in RAID group %s%s: %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group

**reason** (STRING): The reason the reconstruction cannot start



## raid.rg.recons.cksum.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch during the RAID reconstruction process. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry. This indicates that the block is read from the wrong location. Data ONTAP® takes appropriate recovery actions, as described in other logged events.

### Corrective Action

(None).

### Syslog Message

There was a block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during reconstruction.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier ID.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during the RAID reconstruction process. The embedded checksum computed over the checksum entry does not match the stored embedded checksum in the checksum entry, indicating the corruption of the checksum entry. Data ONTAP® takes appropriate recovery actions, as described in other logged events.

### Corrective Action

(None).

## Syslog Message

Invalid checksum entry on %s%s, block #%llu, during reconstruction.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier ID, and site.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.cksum.err

## Severity

NOTICE

## Description

This message occurs when the system detects a checksum error on the stripe during a RAID reconstruction operation. The checksum computed is not the same as the stored checksum. Data ONTAP® takes appropriate recovery actions, as described in additional logged events

## Corrective Action

(None).

## Syslog Message

Checksum error on %s%s, block #%llu during reconstruction.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier ID, and site.  
**blockNum** (LONGINT): Physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.cksum.rderr

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum block media error in an advanced\_zoned checksum's (AZCS) RAID group during RAID reconstruction. Data ONTAP® takes appropriate recovery actions, as described in other logged events.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [%llu - %llu] during reconstruction.

### Parameters

**owner** (STRING): Owner of the disk.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier ID.

**blockNum** (LONGINT): First physical disk block number containing the error.

**LblockNum** (LONGINT): Last physical disk block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during the RAID reconstruction process. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® takes appropriate recovery actions, as described in other logged events.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block %llu, during reconstruction.

### Parameters

**owner** (STRING): Owner of the disk.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier ID.

**blockNum** (LONGINT): Physical disk block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.deferred

### Severity

NOTICE

### Description

This event is issued when the a raid group reconstruction has been deferred due to inadequate incore resources.

### Corrective Action

(None).

### Syslog Message

%s%s: reconstruction deferred (%s)

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is unable to reconstruct

**reason** (STRING): A text reason code

## raid.rg.recons.deferred.ok

### Severity

NOTICE

### Description

A previously deferred reconstruction is now ready to proceed.

### Corrective Action

(None).

### Syslog Message

%s%s: reconstruction previously deferred, now proceeding

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is now able to reconstruct

## raid.rg.recons.disabled

### Severity

ERROR

### Description

This message occurs when a RAID group needs to start reconstruction, but reconstruction is disabled for the RAID group.

### Corrective Action

Run the 'aggr recons enable' command for the RAID group.

### Syslog Message

Reconstruction needs to be performed on RAID group %s%s, but reconstruction is disabled.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): RAID group name.

## raid.rg.recons.done

### Severity

NOTICE

### Description

This message occurs when reconstruction has been completed on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: reconstruction for %s completed in %s.

### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.

**owner** (STRING): Owner of the raid group.

**disk** (STRING): Name of the disk object for which reconstruction completed.

**duration** (STRING): The amount of time the reconstruction required

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.recons.info

### Severity

NOTICE

### Description

This message occurs when a disk is selected for RAID group reconstruction.

### Corrective Action

(None).

### Syslog Message

Spare disk %s will be used to reconstruct one missing disk in RAID group %s%s.

### Parameters

**disk** (STRING): Name of the spare disk object that will be used for reconstruction.

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

## raid.rg.recons.media.err.bypass

### Severity

ERROR

### Description

This event is issued when reconstruction detects a read error but the error is bypassed. The block with read error is then zeroed and written to disk.

### Corrective Action

Ensure wafliion successfully completes. Wafliion will preserve the filesystem consistency.

### Syslog Message

Media error bypass enabled, zeroing block %llu of Disk %s%s

### Parameters

**blockNum** (LONGINT): The physical block number that hit the read error

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk that hit the read error

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.recons.mismatch

### Severity

ERROR

### Description

There was no available disk for reconstruction matching all parameters of an ideal replacement in the RAID group. Data ONTAP selected a disk with different interface speed or rpm, or from the opposite spare pool in a sync-mirror configuration.

### Corrective Action

Wait for reconstruction to complete. Install an appropriate spare disk. Fail the mismatched disk using 'disk fail'. It will be copied to the new spare disk. After copying is done, use 'disk unfail' to unfail the mismatched disk and convert it to a spare.

### Syslog Message

%s%s: matching disk not available for reconstruction; using disk %s %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group

**disk** (STRING): The name of the disk used for reconstruction

**detail** (STRING): Description of disk mismatch; one or more of: from the opposite spare pool, with different rpm.

## raid.rg.recons.missing

### Severity

NOTICE

### Description

Information about disks missing from the RAID group.

### Corrective Action

(None).

### Syslog Message

RAID group %s%s is missing %d disk(s).

### Parameters

**owner** (STRING): The owner of the RAID group

**rg** (STRING): The name of the RAID group

**missing** (INT): Number of disks missing (1 or 2)

## raid.rg.recons.multidisk

### Severity

NOTICE

### Description

This event is issued when the reconstruction encounters an unrecoverable I/O error. We try to avoid panic in this case.

**Corrective Action**

(None).

**Syslog Message**

%s%s: reconstruction encountered a multi-disk error

**Parameters**

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the reconstructing raid group object.

**raid.rg.recons.multidisk.bypass****Severity**

NOTICE

**Description**

Reconstruction can proceed and it will ignore media errors until completion.

**Corrective Action**

(None).

**Syslog Message**

%s%s: reconstruction will bypass media errors

**Parameters**

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the reconstructing raid group object.

**raid.rg.recons.multidisk.fail****Severity**

ALERT

**Description**

The recovery procedure for a media error during reconstruction has failed. The filer will panic to allow recovery after reboot. After reboot, the affected volume will be restricted and marked waf1 inconsistent.

**Corrective Action**

Start waf1ron to bring the affected volume online.

**Syslog Message**

%s%s: failed to recover from a multi-disk error

**Parameters**

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the reconstructing raid group object.

**raid.rg.recons.resume**



## Severity

NOTICE

## Description

This message occurs when reconstruction has been resumed on a specific raid group.

## Corrective Action

(None).

## Syslog Message

%s%s: resuming reconstruction, using disk %s (disk block %llu, %d%% complete)

## Parameters

**owner** (STRING): Owner of the raid group.

**rg** (STRING): Name of the raid group object that is initiating reconstruction.

**disk** (STRING): Name of the disk object that is the reconstruction target.

**blockNum** (LONGINT): The physical block number on which reconstruction is resuming.

**percent** (INT): The percentage of reconstruction already complete.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.recons.start

## Severity

NOTICE

## Description

This message occurs when reconstruction has been initiated on a specific raid group.

## Corrective Action

(None).

## Syslog Message

%s%s: starting reconstruction, using disk %s, disk block %llu.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, carrier serial number, and site.

**owner** (STRING): Owner of the raid group.

**disk** (STRING): Name of the disk object that is the reconstruction target

**startBlockNum** (LONGINT): The physical block number on which reconstruction is starting.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.recons.suspended

### Severity

NOTICE

### Description

This message occurs when reconstruction is suspended on a specific RAID group that is yielding to higher priority reconstruction.

### Corrective Action

(None).

### Syslog Message

%s%s: reconstruction suspended at disk block %llu after %s, yielding to higher priority reconstruction

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group object that suspended reconstruction.

**blockNum** (LONGINT): Physical disk block number where reconstruction stopped.

**duration** (STRING): Amount of time spent on reconstruction before it was suspended.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.recons.throttle

### Severity

NOTICE

### Description

This event is issued when we start throttling the number of messages that reconstruction generates due to an overwhelming number of messages that are being generated.

### Corrective Action

(None).

### Syslog Message

Too many bad checksum messages; will emit message once every %d seconds, dropping others.

### Parameters

**interval** (INT): The interval at which we will generate messages

## raid.rg.reparity.aborted

### Severity

NOTICE

### Description

This event is issued when parity recomputation has been aborted on a specific raid group.

### Corrective Action

(None).

## Syslog Message

%s: parity recomputation aborted after %s

## Parameters

**rg** (STRING): The name of the raid group object that is completed parity recomputation

**duration** (STRING): The amount of time the reparity ran

## raid.rg.reparity.bad.cksum

### Severity

NOTICE

### Description

This message occurs when the computed checksum of a block does not match the stored checksum during a reparity operation.

### Corrective Action

(None).

## Syslog Message

Checksum mismatch on %s%s, block #%llu during parity recomputation; data is assumed to be correct and fixing checksum.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical block number containing the error.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.reparity.cksum.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch during parity recomputation. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).

## Syslog Message

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during parity recomputation.

## Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.

**expected\_dbn** (LONGINT): Expected physical disk block number.

**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.

**expected\_vbn** (LONGINT): Expected volume block number.

**shelf** (STRING): Disk Shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.reparity.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during parity recomputation. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s%s, block #%llu, during parity recomputation.

### Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): Physical disk block number containing the error.

**shelf** (STRING): Disk shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.reparity.cksum.rderr

### Severity

NOTICE

### Description

This message occurs when parity recomputation detects a checksum block media error in an advanced\_zoned checksum's (azcs) RAID group. Data ONTAP® takes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during parity recomputation.

### Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information about the disk object that contains the error.

**blockNum** (LONGINT): First physical disk block number containing the error.

**LblockNum** (LONGINT): Last physical disk block number containing the error.

**shelf** (STRING): Disk shelf identifier for the disk shelf on which the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf on which the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.reparity.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during parity recomputation. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

## Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during parity recomputation.

## Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): Physical disk block number containing the error.

**shelf** (STRING): Disk Shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the shelf where disk is located.

**vendor** (STRING): Name of the disk drive vendor.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.reparity.deferred

### Severity

NOTICE

### Description

This event is issued when the a raid group reparity has been deferred due to inadequate incore resources.

### Corrective Action

(None).

### Syslog Message

%s: parity recomputation deferred (%s)

### Parameters

**rg** (STRING): The name of the raid group object that is unable to reparity

**reason** (STRING): A text reason code

## raid.rg.reparity.deferred.ok

### Severity

INFORMATIONAL

### Description

This message occurs when a previously deferred reparity operation is now ready to proceed because of availability of incore resources.

### Corrective Action

(None).

### Syslog Message

%s: parity recomputation previously deferred is now proceeding

## Parameters

**rg** (STRING): Name of the RAID group object that can proceed with reparity operation.

## raid.rg.reparity.done

### Severity

NOTICE

### Description

This event is issued when parity recomputation has been completed on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s: parity recomputation completed in %s

### Parameters

**rg** (STRING): The name of the raid group object that is completed parity recomputation

**duration** (STRING): The amount of time the reparity required

## raid.rg.reparity.media.err

### Severity

NOTICE

### Description

This event is issued when reparity i/o detects a medium error

### Corrective Action

(None).

### Syslog Message

Read error on %s%s, block #%llu during parity recomputation

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## **raid.rg.reparity.rewrite**

### **Severity**

NOTICE

### **Description**

This event is issued when we rewrite parity on a stripe.

### **Corrective Action**

(None).

### **Syslog Message**

Rewriting parity on RAID group %s, stripe #%llu.

### **Parameters**

**rg** (STRING): The name of the RAID group

**stripe** (LONGINT): The stripe number

## **raid.rg.reparity.start**

### **Severity**

NOTICE

### **Description**

This event is issued when parity recomputation has been initiated on a specific raid group.

### **Corrective Action**

(None).

### **Syslog Message**

%s: starting parity recomputation

### **Parameters**

**rg** (STRING): The name of the raid group object that is initiating parity recomputation

## **raid.rg.reparity.throttle**

### **Severity**

NOTICE

### **Description**

This event is issued when we start throttling the number of messages that reparity generates due to an overwhelming number of messages that are being generated.

### **Corrective Action**

(None).

### **Syslog Message**

Too many parity/bad checksum messages; will emit message once every %d seconds, dropping others.



## Parameters

**interval** (INT): The interval at which we will generate messages

## raid.rg.resync.aborted

### Severity

NOTICE

### Description

This event is issued when resynchronization has been aborted on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: resynchronization aborted after %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is aborting resynchronization

**duration** (STRING): The amount of time the resynchronization ran

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.resync.done

### Severity

NOTICE

### Description

This event is issued when resynchronization has been completed on a specific raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: resynchronization completed in %s

### Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is initiating resynchronization

**duration** (STRING): The amount of time the resynchronization required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.resync.info

### Severity

NOTICE

### Description

This message occurs periodically at the end of RAID group resynchronization, providing information such as the number of reads and writes issued, and number of blocks copied.

## Corrective Action

(None).

## Syslog Message

"%s": number of reads:"%ld" number of writes:"%ld" number of parity blocks read:"%ld" number of blocks copied:"%ld" at throughput rate: %d MB/s.

## Parameters

**rg** (STRING): RAID group in the destination plex.  
**nreadio** (LONGINT): Number of disk reads issued.  
**nwriteio** (LONGINT): Number of tetris writes issued.  
**nparity\_blocks\_read** (LONGINT): Number of parity blocks read.  
**nblocks\_copied** (LONGINT): Number of blocks copied.  
**throughput** (INT): Current throughput rate.

## raid.rg.resync.snap.list

### Severity

INFORMATIONAL

### Description

This message occurs at the beginning and at the end of a resynchronization and it lists a bitmap of Snapshot(tm) copies involved in the resynchronization.

### Corrective Action

(None).

### Syslog Message

%s %s%s: resync Snapshot copies %x %x %x %x %x %x %x %x

### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**rg** (STRING): RAID group in the resynchronizing plex.  
**snapmap7** (INT): Bitmap for Snapshot copies 255-224.  
**snapmap6** (INT): Bitmap for Snapshot copies 223-192.  
**snapmap5** (INT): Bitmap for Snapshot copies 191-160.  
**snapmap4** (INT): Bitmap for Snapshot copies 159-128.  
**snapmap3** (INT): Bitmap for Snapshot copies 127-96.  
**snapmap2** (INT): Bitmap for Snapshot copies 95-64.  
**snapmap1** (INT): Bitmap for Snapshot copies 63-32.  
**snapmap0** (INT): Bitmap for Snapshot copies 31-0.

## raid.rg.resync.start

### Severity

NOTICE

### Description

This event is issued when resynchronization has been initiated on a specific raid group.

## Corrective Action

(None).

## Syslog Message

%s%s: start resynchronization (level %d)

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is initiating resynchronization

**level** (INT): (None).

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.sameCarrier

### Severity

NOTICE

### Description

This message occurs when the disks selected during aggregate creation and mirroring, disk addition, disk replacement, or reconstruction results in two disks from the same disk carrier being part of a single RAID group. The carrier should not be removed; otherwise the RAID group is doubly degraded, which significantly reduces performance. This situation also increases the chance of data loss. Data loss could occur because of either the user forcing the command to make such a disk selection, or Data ONTAP® could not find more suitable disks. This event is generated for every pair of disks residing in a disk carrier that would become part of a single RAID group. Data ONTAP makes appropriate recovery actions. One or more disk copies are performed to correct the problem.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk1** (STRING): Name of the first disk in the disk carrier.

**disk2** (STRING): Name of the second disk in the disk carrier.

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

## raid.rg.scrub.cksum.bno

### Severity

NOTICE

### Description

This message occurs when the system detects a block number mismatch during RAID parity scrub. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry, indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).

## Syslog Message

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during RAID parity scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during RAID parity scrub. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

## Corrective Action

(None).

## Syslog Message

Invalid checksum entry on %s%s, block #%llu, during RAID parity scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Disk Shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.cksum.err

### Severity

NOTICE

### Description

This event is issued when scrubbing detects a single-disk checksum error on a stripe.

### Corrective Action

(None).

### Syslog Message

Checksum error on %s%s, block #%llu during scrub

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.cksum.rderr

### Severity

NOTICE

### Description

This message occurs when background scrub detects a single-disk checksum block media error in an advanced\_zoned checksum's (azcs) RAID group. Data ONTAP® takes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): First physical disk block number containing the error.  
**LblockNum** (LONGINT): Last physical disk block number containing the error.  
**shelf** (STRING): Disk shelf identifier for the disk shelf on which the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf on which the disk drive is located.  
**vendor** (STRING): Name of the vendor of the disk drive.  
**model** (STRING): Model of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during RAID parity scrub. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during RAID parity scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Disk Shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.context.valid

**Severity**

NOTICE

**Description**

This message occurs when scrubbing is completed on a specific RAID group and its context now has parity protection.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group object that completed the scrub.

**vol\_context** (STRING): Status of context parity protection of the aggregate.

**raid.rg.scrub.cxt.restart****Severity**

NOTICE

**Description**

This message occurs when a suspended regular scrub is restarting as a context scrub from the beginning instead of resuming regular scrub at the last suspended point. It can happen when a RAID group requires parity updating for context parity protection.

**Corrective Action**

(None).

**Syslog Message**

%s%s: restarting suspended scrub for WAFL context parity protection.

**Parameters**

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

**raid.rg.scrub.deferred****Severity**

NOTICE

**Description**

This event is issued when the a raid group scrub has been deferred due to inadequate incore resources.

**Corrective Action**

(None).

## Syslog Message

%s%s: scrub deferred (%s)

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is unable to scrub

**reason** (STRING): A text reason code

## raid.rg.scrub.deferred.ok

### Severity

NOTICE

### Description

A previously deferred scrub is now ready to proceed.

### Corrective Action

(None).

## Syslog Message

%s%s: scrub previously deferred, now proceeding

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is now able to scrub

## raid.rg.scrub.done

### Severity

NOTICE

### Description

This event is issued when scrubbing has been completed on a specific raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: scrub completed in %s

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is completed scrub

**duration** (STRING): The amount of time the scrubbing required

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.scrub.lw.aborted

### Severity

NOTICE



## Description

This message occurs when the system aborts a special type of RAID scrub because it detected an inconsistent RAID write signature on the specified RAID group. The RAID scrub is required before lost-write protection for RAID I/O can be used on the RAID group.

## Corrective Action

(None).

## Syslog Message

%s%s stripe #%llu: scrub needed to enable lost-write protection for RAID I/O aborted after %s.

## Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group object for which the lost-write scrub was aborted.

**stripe** (LONGINT): RAID stripe on which the inconsistent write signature was detected.

**duration** (STRING): Amount of time the lost-write scrub has taken.

## raid.rg.scrub.lw.restart

### Severity

NOTICE

### Description

This message occurs when a suspended regular scrub is restarting from the beginning to allow lost-write protection for RAID I/O instead of resuming regular scrub at the last suspended point. This happens when lost-write protection for RAID I/O is enabled on this RAID group while a regular scrub is suspended.

### Corrective Action

(None).

### Syslog Message

%s%s: restarting suspended scrub to enable lost-write protection for RAID I/O.

### Parameters

**owner** (STRING): Owner of the RAID group.

**rg** (STRING): Name of the RAID group.

## raid.rg.scrub.media.err

### Severity

NOTICE

### Description

This message occurs when scrub I/O detects a medium error on a disk that is also called as an unrecoverable access error. This error is fixed immediately by reconstructing blocks from parity data.

### Corrective Action

(None).

### Syslog Message

Read error on %s%s, block #%llu, during scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Name of the disk object that contains the error.  
**blockNum** (LONGINT): Physical block number containing the error.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.media.err.bypass

### Severity

INFORMATIONAL

### Description

This event is issued when a read error is detected but ignored due to the presence of media bypass (ignore double errors).

### Corrective Action

(None).

### Syslog Message

Media error bypass enabled, zeroing disk block %llu of %s%s

## Parameters

**blockNum** (LONGINT): The physical block number containing the error.  
**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.media.errs

### Severity

NOTICE

## Description

This event is issued when scrubbing detects a multi-disk media error on a stripe.

## Corrective Action

(None).

## Syslog Message

Multiple bad blocks on %s%s, block #%llu during scrub

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): The name of the raid group object that contains the multi-disk error.

**blockNum** (LONGINT): The physical block number containing the error.

## raid.rg.scrub.media.recommend.reassign.err

### Severity

NOTICE

### Description

This event is issued when scrub i/o detects a recommend reassign error

### Corrective Action

(None).

### Syslog Message

Block recommended for reassignment on %s%s, block #%llu during scrub

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk object that contains the error.

**blockNum** (LONGINT): The physical block number containing the error.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.pi

### Severity

NOTICE

### Description

This message occurs when RAID scrub finds an inconsistent parity stripe. Data ONTAP® makes appropriate recovery actions, as described in other logged events.

**Corrective Action**

(None).

**Syslog Message**

Inconsistent parity on RAID group %s, stripe #%llu.

**Parameters**

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

**raid.rg.scrub.resume****Severity**

NOTICE

**Description**

This event is issued when scrubbing has been resumed on a specific raid group.

**Corrective Action**

(None).

**Syslog Message**

%s%s: resuming scrub at stripe %llu (%d%% complete)

**Parameters**

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is resuming scrubbing

**stripe** (LONGINT): resuming at stripe.

**percentage** (INT): (None).

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**raid.rg.scrub.rewrite.data****Severity**

NOTICE

**Description**

This message occurs when scrub fixes a parity inconsistency error.

**Corrective Action**

(None).

**Syslog Message**

Scrub fixing bad data on disk %s%s, block #%llu.

**Parameters**

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Formatted information of the disk drive containing the block being rewritten.

**blockNum** (LONGINT): Disk block number.

**shelf** (STRING): Disk shelf identifier where the disk drive is located.

**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.

**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.rewrite.parity

### Severity

NOTICE

### Description

This message occurs when scrub fixes a parity inconsistency error.

### Corrective Action

(None).

### Syslog Message

Scrub fixing bad parity on disk %s%s, block #%llu.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk drive containing the block being rewritten.  
**blockNum** (LONGINT): Disk block number.  
**shelf** (STRING): Disk shelf identifier where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where disk drive is located.  
**vendor** (STRING): Name of the vendor of the disk drive.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.rg.scrub.start

### Severity

NOTICE

### Description

This event is issued when scrubbing has been initiated on a specific raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: starting scrub

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is initiating scrubbing

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.scrub.start.zone

### Severity

NOTICE

### Description

This event is issued when scrubbing has been initiated on a specific raid group for a specific range.

### Corrective Action

(None).

## Syslog Message

%s%s: starting scrub at range [%llu - %llu]

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that is initiating scrubbing

**begin** (LONGINT): starting at stripe.

**end** (LONGINT): ending at stripe.

## raid.rg.scrub.stopped

### Severity

NOTICE

### Description

This event is issued when scrub is stopped on specified raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: scrub stopped at stripe %llu after %s

## Parameters

**owner** (STRING): The owner of the raid group

**rg** (STRING): The name of the raid group object that the scrub is stopped

**stripe** (LONGINT): stopped at stripe.

**duration** (STRING): The amount of time the scrub has taken.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.rg.scrub.summary.cksum

### Severity

NOTICE

### Description

This event is generated as a summary report for checksum errors found during a scrub.

### Corrective Action

(None).

### Syslog Message

Scrub found %d checksum error%s in %s%s

### Parameters

**errors** (INT): The total number of checksum errors found

**suffix** (STRING): The suffix used to report in syslog

**rg** (STRING): The name of the raid group

**current** (STRING): The current status

## raid.rg.scrub.summary.lw

### Severity

NOTICE

### Description

This message occurs as a summary report for lost-write errors found during a scrub.

### Corrective Action

(None).

### Syslog Message

Scrub found %d RAID write signature inconsistenc%s in %s%s

### Parameters

**errors** (INT): Total number of RAID write signature inconsistencies found.

**suffix** (STRING): Suffix used to report in syslog.

**rg** (STRING): Name of the RAID group.

**current** (STRING): Current status.

## raid.rg.scrub.summary.media

### Severity

NOTICE

### Description

This event is generated as a summary report for media errors found during a scrub.

### Corrective Action

(None).

## Syslog Message

Scrub found %d media error%s in %s%s

## Parameters

**errors** (INT): The total number of media errors found  
**suffix** (STRING): The suffix used to report in syslog  
**rg** (STRING): The name of the raid group  
**current** (STRING): The current status

## raid.rg.scrub.summary.pi

### Severity

NOTICE

### Description

This event is generated as a summary report for parity inconsistencies found during a scrub.

### Corrective Action

(None).

## Syslog Message

Scrub found %d parity inconsistenc%s in %s%s

## Parameters

**pis** (INT): The total number of parity inconsistencies found  
**suffix** (STRING): The suffix used to report in syslog  
**rg** (STRING): The name of the raid group  
**current** (STRING): The current status

## raid.rg.scrub.suspended

### Severity

NOTICE

### Description

This event is issued when scrub is suspended on specified raid group.

### Corrective Action

(None).

## Syslog Message

%s%s: scrub suspended at stripe %llu after %s

## Parameters

**owner** (STRING): The owner of the raid group  
**rg** (STRING): The name of the raid group object that the scrub is suspended  
**stripe** (LONGINT): suspended at stripe.  
**duration** (STRING): The amount of time the scrub has taken.  
**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.



## raid.rg.spares.low

### Severity

ERROR

### Description

This message occurs when the number of available spares for a RAID group drops below the value specified in `raid.min_spare_count`. If a RAID reconstruct operation is required (for example, due to a future disk failure), the system might enter a degraded mode.

### Corrective Action

Assign unowned disks to this system. If there are no unowned disks, physically add more spare disks and assign them to this system.

### Syslog Message

%s

### Parameters

`msg` (STRING): Name of the RAID group.

## raid.rg.trecons

### Severity

NOTICE

### Description

This message occurs when triple reconstruction is initiated after the failure of three disks in RAID-TEC raid group.

### Corrective Action

(None).

### Syslog Message

%s%s: starting triple reconstruction, using disks %s and %s and %s

### Parameters

`owner` (STRING): Owner of the raid group.

`rg` (STRING): Name of the raid group object that is initiating triple reconstruction.

`disk` (STRING): Name of the disk object that is the first reconstruction target.

`disk1` (STRING): Name of the disk object that is the second reconstruction target.

`disk2` (STRING): Name of the disk object that is the third reconstruction target.

`aggregate_uuid` (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.root events

### raid.root.unmirrored

### Severity

ALERT

## Description

This message occurs when the root volume is not mirrored in a MetroCluster(tm) configuration.

## Corrective Action

Mirror the root volume if this storage controller is part of a MetroCluster configuration by using the "storage aggregate mirror" command.

## Syslog Message

The root volume is not mirrored. A takeover/switchover of this storage controller might not be possible if a disaster occurs.

## Parameters

(None).

## raid.root.vol.noFaultIsolation

### Severity

ALERT

### Description

This message occurs when creation of the initial root aggregate is forced to use disks that would be from both Pool0 and Pool1 if it was to be upgraded to a mirrored root.

### Corrective Action

a) Determine how this occurred: possible cause is a wiring problem. b) Based on the information you gather, determine how to correct the issue. For example, use the "storage disk replace" command to copy the disk belonging to the wrong pool to a disk belonging to the right pool. For MetroCluster configurations, this would mean using "storage disk replace" to copy any pool1 disk in the root to a pool0 disk. c) If you need assistance, contact NetApp technical support.

### Syslog Message

Root volume %s%s: Current plex will have disks from mixed pools if the root volume is upgraded to mirror.

### Parameters

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.scrub events

### raid.scrub.failed.to.schedule

#### Severity

NOTICE

#### Description

This message occurs when scrub cannot be scheduled as specified by the "raid.scrub.schedule" option.

#### Corrective Action

(None).

## Syslog Message

Failed to schedule scrub as specified by the "raid.scrub.schedule" option. Default scrub schedule is added.

## Parameters

(None).

## raid.scrub.lw.id.inconsist

### Severity

NOTICE

### Description

This message occurs when the system detects an inconsistent RAID write signature on a RAID stripe during a RAID scrub operation. This inconsistent RAID write signature might be due to a write operation not being successfully committed to disk. Data ONTAP® attempts to recover the bad block automatically. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

### Corrective Action

(None).

## Syslog Message

Inconsistent RAID write signature on RAID group %s%s, stripe #%llu, stripe bitmap %llu, during scrub.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

**more\_stripes** (LONGINT): Bitmap representing all of the stripes that encountered the lost write error. The bit set in this bitmap gives the stripe number relative to the stripes parameter.

## raid.scrub.suspended

### Severity

NOTICE

### Description

This event is issued if automatic/scheduled scrub is suspended on some raid group because of a group state transition.

### Corrective Action

(None).

## Syslog Message

Disk scrub suspended.

## Parameters

(None).

## raid.scrub.suspended.timer

### Severity

NOTICE

### Description

This event is issued if automatic/scheduled scrub is suspended on some raid group because the allocated time limit was exceeded.

### Corrective Action

(None).

### Syslog Message

Disk scrub suspended because the scrub time limit %d was exceeded. It will resume at the next weekly/scheduled scrub.

### Parameters

**duration** (INT): Duration (in mins) after which scrub is automatically suspended.

## raid.shared events

### raid.shared.disk.awaiting.done

### Severity

INFORMATIONAL

### Description

This message occurs when a node receives an indication from its HA partner node that a shared disk is undergoing a change in state and the partner is waiting for the local node to complete a shared state transition.

### Corrective Action

(None).

### Syslog Message

Received shared disk awaiting done %s, state %s, substate 0x%IX, partner state %s, partner substate 0x%IX, partner dblade ID %s host type %d add details %s

### Parameters

**disk\_info** (STRING): Formatted information of the shared disk.

**local\_state** (STRING): State of the shared disk on the local node.

**local\_substate** (LONGINTEX): Sub-state of the shared disk on the local node.

**partner\_state** (STRING): State of the shared disk on the partner node.

**partner\_substate** (LONGINTEX): Sub-state of the shared disk on the partner node.

**received\_ack** (STRING): Dblade ID of node sending await done ack.

**host\_type** (INT): Type of host node sending await done ack.

**add\_details** (STRING): Additional details of the awaiting done event.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.shared.disk.exchange

### Severity

INFORMATIONAL

### Description

This message occurs when a node receives an indication from its HA partner node that a shared disk is undergoing a change in state (e.g., disk is failing, going offline, or is getting unfailed). Accompanying EMS messages will indicate if any corrective action is required.

### Corrective Action

(None).

### Syslog Message

Received shared disk state exchange %s, event %s, state %s, substate 0x%IX, partner state %s, partner substate 0x%IX, failure reason %s, sick reason %s, offline reason %s, online reason %s, partner dblade ID %s, host %d %s

### Parameters

**disk\_info** (STRING): Formatted information of the shared disk.

**event** (STRING): Describes the event which caused the partner to send this state exchange message. If this is a periodic exchange instead of an event triggered exchange, event is NONE.

**local\_state** (STRING): State of the shared disk on the local node.

**local\_substate** (LONGINTEX): Sub-state of the shared disk on the local node.

**partner\_state** (STRING): State of the shared disk on the partner node.

**partner\_substate** (LONGINTEX): Sub-state of the shared disk on the partner node.

**failure\_reason** (STRING): Describes the error leading to the partner node failing the shared disk.

**sick\_reason** (STRING): Describes the error leading to the partner node pre-failing the shared disk.

**offline\_reason** (STRING): Describes the reason the partner node is offlining the shared disk.

**online\_reason** (STRING): Describes the reason the partner node is onlining the shared disk.

**recv\_reply** (STRING): Describes the dblade ID of a partner node that sent a reply.

**host\_type** (INT): Type of host sending the reply.

**add\_details** (STRING): Additional partner node details of the shared disk state exchange.

**timestamp** (INT): Timestamp of the last on-disk state change.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.sp events

## raid.sp.add.aggr.grow.noDone

### Severity

NOTICE

### Description

This message occurs when one or more aggregate's capacity is not increased as part of the storage pool add job.

### Corrective Action

Use the (privilege: diagnostic) "storage pool add -storage-pool <storage pool name> -auto-grow -aggregates" command to retry the operation. If the condition persists, contact NetApp technical support.

### Syslog Message

Could not grow one or more aggregates provisioned from storage pool %s.

### Parameters

**name** (STRING): Name of the storage pool.

**uuid** (STRING): UUID of the storage pool.

**jobid** (LONGINT): ID of the job.

## raid.sp.add.success

### Severity

NOTICE

### Description

This message occurs when disks are added to the storage pool and all the aggregates provisioned from the storage pool have been grown to take capacity from new disks.

### Corrective Action

(None).

### Syslog Message

Disks %s added to storage pool %s.

### Parameters

**disks** (STRING): List of disks added to the storage pool.

**name** (STRING): Name of the storage pool.

**uuid** (STRING): UUID of the storage pool.

**jobid** (LONGINT): ID of the job.

## raid.sp.aggr.add.rollback

### Severity

NOTICE

## Description

This message occurs when available capacity is reassigned to a storage pool because adding storage pool capacity to an aggregate failed. This can happen if the node owning the aggregate is rebooted while additional storage pool capacity is zeroing. Other events provide additional failure information.

## Corrective Action

(None).

## Syslog Message

Aggregate add job for aggregate %s successfully rolled back available capacity to storage pool %s.

## Parameters

**aggr** (STRING): Name of the aggregate to which capacity is being added.

**name** (STRING): Name of the storage pool.

**uuid** (STRING): UUID of the storage pool.

**aggrUuid** (STRING): UUID of the aggregate.

**owner** (STRING): Owning node of the aggregate.

**curUnit** (INT): Current available allocation unit.

**prevUnit** (INT): Allocation unit available before the job started.

**jobid** (LONGINT): ID of the job.

## raid.sp.aggr.add.success

### Severity

NOTICE

### Description

This message occurs when capacity from a storage pool is added to an aggregate.

### Corrective Action

(None).

### Syslog Message

Added %d allocation units from storage pool %s to aggregate %s.

### Parameters

**allocationUnits** (INT): Number of allocation units added.

**name** (STRING): Name of the storage pool.

**aggr** (STRING): Name of the aggregate to which capacity is added.

**uuid** (STRING): UUID of the storage pool.

**aggrUuid** (STRING): UUID of the aggregate.

**jobid** (LONGINT): ID of the job.

## raid.sp.aggr.grow.success

### Severity

NOTICE

### Description

This message occurs when an aggregate capacity is successfully increased to span all disks belonging to a storage pool.

## Corrective Action

(None).

## Syslog Message

%s added to aggregate %s from storage pool %s.

## Parameters

**capacity** (STRING): Amount of capacity added to the aggregate.  
**aggr** (STRING): Name of the aggregate whose capacity is increased.  
**name** (STRING): Name of the storage pool.  
**uuid** (STRING): UUID of the storage pool.  
**aggrUuid** (STRING): UUID of the aggregate.  
**jobid** (LONGINT): ID of the job.

## raid.sp.create.failure

### Severity

ERROR

### Description

This message occurs when a storage pool creation operation fails.

### Corrective Action

Use the "storage pool delete -storage-pool <storage pool name>" command to delete the partially created storage pool and then retry the create operation. If the condition persists, contact NetApp technical support.

### Syslog Message

Storage pool %s was not successfully created. Reason: %s

### Parameters

**name** (STRING): Name of the storage pool.  
**reason** (STRING): Reason for failure.  
**uuid** (STRING): UUID of the storage pool.  
**nodes** (STRING): List of nodes sharing the storage pool.  
**disks** (STRING): List of disks used to create the storage pool.  
**jobid** (LONGINT): ID of the job.

## raid.sp.create.success

### Severity

NOTICE

### Description

This message occurs when a storage pool is created successfully.

### Corrective Action

(None).

### Syslog Message

Storage pool %s was created.



## Parameters

**name** (STRING): Name of the storage pool.  
**uuid** (STRING): UUID of the storage pool.  
**nodes** (STRING): List of nodes sharing the storage pool.  
**disks** (STRING): List of disks used to create the storage pool.  
**jobid** (LONGINT): ID of the job.

## raid.sp.delete.success

### Severity

NOTICE

### Description

This message occurs when a storage pool is deleted successfully.

### Corrective Action

(None).

### Syslog Message

Storage pool %s deleted.

## Parameters

**name** (STRING): Name of the storage pool.  
**uuid** (STRING): UUID of the storage pool.  
**jobid** (LONGINT): ID of the job.

## raid.sp.healthy

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® detects that a storage pool that was unhealthy earlier has become healthy again.

### Corrective Action

(None).

### Syslog Message

Storage pool %s is now healthy.

## Parameters

**name** (STRING): Name of the storage pool.  
**uuid** (STRING): UUID of the storage pool.

## raid.sp.reassign.success

### Severity

NOTICE

## Description

This message occurs when available capacity is reassigned successfully from one node to another node in a storage pool.

## Corrective Action

(None).

## Syslog Message

Reassigned %d allocation units from node %s to node %s for storage pool %s.

## Parameters

**allocationUnits** (INT): Number of allocation units reassigned.  
**source** (STRING): Node that owns the available capacity to be reassigned.  
**destination** (STRING): Node to which the available capacity is reassigned.  
**name** (STRING): Name of the storage pool.  
**uuid** (STRING): UUID of the storage pool.  
**jobid** (LONGINT): ID of the job.

## raid.sp.unhealthy

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® detects that a storage pool has become unhealthy.

### Corrective Action

(None).

### Syslog Message

Storage pool %s is unhealthy. Reason: %s

### Parameters

**name** (STRING): Name of the storage pool.  
**reason** (STRING): Reason for the storage pool being unhealthy.  
**uuid** (STRING): UUID of the storage pool.

## raid.sp.unlock.failed

### Severity

ALERT

### Description

This message occurs when the storage pool cannot be unlocked after the storage pool operation has finished.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Failed to unlock the storage pool %s. Reason: %s

## Parameters

**name** (STRING): Name of the storage pool.  
**reason** (STRING): Reason for failure.  
**uuid** (STRING): UUID of the storage pool.  
**jobid** (LONGINT): ID of the job.

# raid.sparecore events

## raid.sparecore.disk.ownership.change

### Severity

NOTICE

### Description

This message occurs when a sparecore disk is prepared for ownership change by administrator forcibly.

### Corrective Action

(None).

### Syslog Message

Ownership of a sparecore disk %s, is being changed by administator forcibly.

## Parameters

**disk\_info** (STRING): Formatted information of the disk.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.spares events

## raid.spares.media\_scrub.media.err

### Severity

NOTICE

### Description

This event is issued when a media\_scrub i/o detects a medium error on a spare disk

### Corrective Action

(None).

## Syslog Message

Read error on %s%s, block #%llu during media\_scrub

## Parameters

**owner** (STRING): String indicating the owner of the affected spare drive.  
**disk\_info** (STRING): The name of the disk object that contains the error.  
**blockNum** (LONGINT): The physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.spares.media\_scrub.media.recommend.reassign.err

### Severity

NOTICE

### Description

This event is issued when a background media scrub i/o detects a recommend reassign error on a spare drive

### Corrective Action

(None).

## Syslog Message

Block recommended for reassignment on %s%s, block #%llu during media\_scrub

## Parameters

**owner** (STRING): String indicating the owner of the affected spare drive.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): The physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.spares.media\_scrub.rewrite.chunk

### Severity

INFORMATIONAL

### Description

This message occurs when background media scrub rewrites a chunk of blocks for a spare ATA drive.

### Corrective Action

(None).

### Syslog Message

media\_scrub rewriting chunk on %s%s, start block #%llu, count #%d.

### Parameters

**owner** (STRING): Owner of the affected spare disk.

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**blockNum** (LONGINT): Physical block number at which the chunk starts.

**count** (INT): Number of blocks being rewritten.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model name of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.spares.media\_scrub.start

### Severity

NOTICE

### Description

This event is issued when media\_scrub is initiated on a spare drive.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): String indicating the owner of the affected spare drive.

**disk\_info** (STRING): The name of the disk object that contains the error.

**blockNum** (LONGINT): Disk block number at which media scrub started.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.spares.media\_scrub.stopped

### Severity

NOTICE

### Description

This event is issued when media\_scrub stops on a spare drive.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): String indicating the owner of the affected spare drive.  
**disk\_info** (STRING): The name of the disk object that contains the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.spares.media\_scrub.suspend

### Severity

NOTICE

### Description

This event is issued when media\_scrub suspends on a spare drive.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**owner** (STRING): String indicating the owner of the affected spare drive.

**disk\_info** (STRING): The name of the disk object that contains the error.

**blockNum** (LONGINT): Disk block number at which media scrub suspended

**percentage** (INT): Percentage scan completed.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# raid.stripe events

## raid.stripe.replay.bad.block

### Severity

ERROR

### Description

This event is issued when we encounter a bad block buffer entry in NVRAM during stripe replay (indicating corrupted NVRAM). Reparity will start automatically and update parity in stripes involved in recent writes.

### Corrective Action

Need to replace NVRAM if observed repeatedly during different boot-ups.

### Syslog Message

Bad block buffer entry in NVRAM.

### Parameters

(None).

## raid.stripe.replay.bad.checksum

### Severity

ERROR

### Description

This event is issued when we encounter a checksum error in NVRAM during stripe replay (indicating corrupted NVRAM). Reparity will start automatically and update parity in stripes involved in recent writes.

**Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

**Syslog Message**

Checksum error in NVRAM.

**Parameters**

(None).

**raid.stripe.replay.bad.key****Severity**

ERROR

**Description**

This event is issued when we cannot find a block buffer using a key during stripe replay (indicating corrupted NVRAM). Reparity will start automatically and update parity in stripes involved in recent writes.

**Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

**Syslog Message**

Stripe %llu: cannot find NVRAM slot using key %d in %s %s.

**Parameters**

**stripe** (LONGINT): The stripe number.

**key** (INT): The key used.

**vol\_type** (STRING): Volume type.

**vol** (STRING): Aggregate name if found, otherwise NOT FOUND.

**raid.stripe.replay.bad.stripe****Severity**

ERROR

**Description**

This event is issued when we encounter a bad stripe entry in NVRAM during stripe replay (indicating corrupted NVRAM). Reparity will start automatically and update parity in stripes involved in recent writes.

**Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

**Syslog Message**

Bad stripe entry in NVRAM.

**Parameters**

(None).

**raid.stripe.replay.nvgen.mismatch**



**Severity**

NOTICE

**Description**

This message occurs when RAID encounters an NVRAM generation mismatch for an aggregate during stripe replay. Reparity will start automatically and update parity in stripes involved in recent writes.

**Corrective Action**

(None).

**Syslog Message**

NVRAM generation mismatch for %s %s%s.

**Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**raid.stripe.replay.plex.unavailable****Severity**

NOTICE

**Description**

This message occurs when RAID encounters an NVRAM stripe entry for an offline or failed plex during stripe replay. Reparity will start automatically when the plex comes online. At that time, reparity will update parity in stripes involved in recent writes.

**Corrective Action**

(None).

**Syslog Message**

Plex %s in %s %s%s is offline or failed.

**Parameters**

**plex** (STRING): Name of the plex.

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**raid.stripe.replay.pzero.mismatch****Severity**

ERROR

**Description**

This event is issued when the checksum of the parity block (as stored in the stripe slot) does not match the actual checksum of the parity block. This indicates corrupted NVRAM contents. Reparity will start automatically and update parity in stripes involved in recent writes.

**Corrective Action**

Need to replace NVRAM if observed repeatedly during different boot-ups.

**Syslog Message**

Parity block checksum mismatch detected in NVRAM for stripe %llu in %s %s.

**Parameters**

**stripe** (LONGINT): The stripe number.

**vol\_type** (STRING): Volume type.

**vol** (STRING): Aggregate name if found, otherwise NOT FOUND.

**raid.stripe.replay.rg.failed****Severity**

NOTICE

**Description**

This message occurs when RAID encounters a RAID group failure during stripe RAID replay. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

**Corrective Action**

(None).

**Syslog Message**

Replay of %s %s%s, plex %s, RAID group %s, stripe %llu failed because of RAID group failure.

**Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**plex** (STRING): Name of the plex.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

**raid.stripe.replay.rg.invalid****Severity**

NOTICE

**Description**

This message occurs when RAID encounters an NVRAM stripe entry for a missing RAID group during stripe replay.

**Corrective Action**

(None).

**Syslog Message**

Invalid RAID group identifier %d in %s %s%s, plex %s.

**Parameters**

**rgid** (INT): Identifier of the RAID group.

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.  
**plex** (STRING): Name of the plex.

## **raid.stripe.replay.summary**

### **Severity**

INFORMATIONAL

### **Description**

This event is generated when stripe replay is complete.

### **Corrective Action**

(None).

### **Syslog Message**

Replayed %u stripes.

### **Parameters**

**nstripes** (INT): The number of stripes

## **raid.stripe.replay.vol.cleantly.shutdown**

### **Severity**

NOTICE

### **Description**

This message occurs when RAID encounters an NVRAM stripe entry for a cleanly shutdown aggregate during stripe replay.

### **Corrective Action**

(None).

### **Syslog Message**

%s %s%s was cleanly shut down.

### **Parameters**

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.

## **raid.stripe.replay.volume.missing**

### **Severity**

NOTICE

### **Description**

This message occurs when RAID encounters an NVRAM stripe entry for a missing aggregate during stripe replay. Reparity will start automatically when the aggregate appears again. At that time, reparity will update parity in stripes involved in recent writes.

**Corrective Action**

(None).

**Syslog Message**

Cannot find aggregate with identifier %s.

**Parameters**

**volId** (STRING): Identifier of the aggregate.

**raid.stripe.replay.volume.offline****Severity**

NOTICE

**Description**

This message occurs when RAID encounters an NVRAM stripe entry for an offline aggregate during stripe replay. Reparity will start automatically when the aggregate comes online. At that time, reparity will update parity in stripes involved in recent writes.

**Corrective Action**

(None).

**Syslog Message**

%s %s %s is not online.

**Parameters**

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the affected aggregate.  
**vol** (STRING): Name of the aggregate.

**raid.switchoverfail events****raid.switchoverFail****Severity**

ERROR

**Description**

This message occurs when a MetroCluster(tm) switchover attempt fails because of an error detected during RAID assimilation.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Switchover assimilation failed: %s.

**Parameters**

**reason** (STRING): Reason for the switchover assimilation failure.

# raid.switchovernodisks events

## raid.switchoverNoDisks

### Severity

INFORMATIONAL

### Description

This message occurs when a MetroCluster(tm) switchover finds no disks during RAID assimilation.

### Corrective Action

(None).

### Syslog Message

MetroCluster switchover assimilation found no disks.

### Parameters

(None).

# raid.syncmirr events

## raid.syncmirr.at.IOrecovery

### Severity

ERROR

### Description

This message occurs in a SyncMirror® configuration (with SyncMirror aggressive timeout capability), when RAID cannot service earlier timed-out I/O from the other plex.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem aggregate, offline plex, and plex mirror status from the event logs. 2. Determine the reason for attempting to bring the plex offline, such as adapter, controller or multidisk failures. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support.

### Syslog Message

Expect longer I/O service time on plex %s due to error status %s.

### Parameters

**plex** (STRING): Name of the plex.

**status** (STRING): Status while the plex was brought offline.

# raid.syncmirror events

## raid.syncmirror.agrsv.to

### Severity

ERROR

## Description

This message occurs in a SyncMirror® configuration (with SyncMirror aggressive timeout capability), when a disk I/O times out with an error.

## Corrective Action

The system will attempt to recover using the other plex. If the error persists, consider replacing the drive or the adapter.

## Syslog Message

Disk driver aggressively timed out an I/O to %s%s belonging to a SyncMirror plex %s. RAID will now attempt to bring the plex offline and redrive the I/O from the other plex.

## Parameters

**owner** (STRING): String indicating the owner of the affected disk.

**disk\_info** (STRING): Name of the disk.

**plex** (STRING): Name of the affected plex.

## raid.syncmirror.at.degraded

### Severity

ALERT

### Description

This message occurs in a SyncMirror® configuration (with SyncMirror aggressive timeout capability), when a mirror goes into degraded state as a result of an offline plex.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem aggregate, the offline plex, and plex mirror status from the event logs. 2. Determine the reason for the offline plex, such as adapter, controller or multidisk failures. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support. 4. Bring the plex online manually after you replace the faulty components.

### Syslog Message

%s%s [UUID:%s] is mirrored and %s plex is offline. It is no longer protected by mirroring.

### Parameters

**owner** (STRING): Owner of the aggregate.

**aggregate** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

**plex** (STRING): Name of the offline plex.

## raid.syncmirror.in.use

### Severity

INFORMATIONAL

### Description

This message occurs when the SyncMirror® license is currently in use.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**nAggrs** (INT): Total number of aggregates on the system.

**nMirrored** (INT): Number of aggregates mirrored.

**rootMirrored** (STRING): Indicates whether root aggregate is mirrored or not.

**raid.syncmirror.plex.online****Severity**

NOTICE

**Description**

This message occurs in a SyncMirror® configuration, when bringing online an offline plex.

**Corrective Action**

(None).

**Syslog Message**

The plex %s is now online.

**Parameters**

**plex** (STRING): Name of the plex.

**raid.tetris events****raid.tetris.cksum.bno****Severity**

NOTICE

**Description**

This message occurs when the system detects a block number mismatch during a write operation. The expected Virtual Block Number (VBN)/Disk Block Number (DBN) is not the same as the stored VBN/DBN from the checksum entry indicating that the block is read from the wrong location. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

**Corrective Action**

(None).

**Syslog Message**

Block number mismatch on %s%s: stored\_dbn = %u, expected\_dbn = %llu; stored\_vbn = %llu, expected\_vbn = %llu during write operation.

**Parameters**

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object that contains the error, including disk name, path,

shelf, bay, serial number, vendor, model, rpm, and carrier serial number.  
**stored\_dbn** (INT): Physical disk block number stored in the checksum entry.  
**expected\_dbn** (LONGINT): Expected physical disk block number.  
**stored\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**expected\_vbn** (LONGINT): Expected volume block number.  
**shelf** (STRING): Identifier of the shelf where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.tetris.cksum.embed

### Severity

NOTICE

### Description

This message occurs when the system detects an invalid checksum entry during a write operation. The embedded checksum computed over the checksum entry do not match, indicating the corruption of the checksum entry. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Invalid checksum entry on %s%s, block #%llu, during write operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object that contains the error, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Identifier of the shelf where the disk drive is located.  
**bay** (STRING): Disk bay within the shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.



## raid.tetris.cksum.err

### Severity

NOTICE

### Description

This message occurs when tetris I/O detects a checksum error. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Checksum error on %s%s, block #%llu during stripe write.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**disk\_info** (STRING): Information about the disk object that contains the error, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.

**blockNum** (LONGINT): Physical block number containing the error.

**shelf** (STRING): Identifier of the shelf where the disk drive is located.

**bay** (STRING): Disk bay within the shelf where the disk drive is located.

**vendor** (STRING): Name of the vendor of the disk drive.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk drive.

**serialno** (STRING): Serial number of the disk drive.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.tetris.cksum.rderr

### Severity

NOTICE

### Description

This message occurs when the system detects a checksum block media error in an advanced\_zoned checksum's (AZCS) RAID group during a stripe write operation. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Checksum block read error on %s%s for blocks [#%llu - #%llu] during stripe write.

### Parameters

**owner** (STRING): Owner of the disk drive.

**disk\_info** (STRING): Information about the disk object that contains the error, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.

**blockNum** (LONGINT): First physical disk block number containing the error.  
**LblockNum** (LONGINT): Last physical disk block number containing the error.  
**shelf** (STRING): Identifier of the shelf where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.tetris.cksum.zero

### Severity

NOTICE

### Description

This message occurs when the system detects an empty checksum entry during a write operation. The checksum entry is zeroed, but the corresponding block is not zeroed. Data ONTAP® makes appropriate recovery actions. Other events describe those actions.

### Corrective Action

(None).

### Syslog Message

Empty checksum entry for non-zeroed block on %s%s, block #%llu, during a write operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object that contains the error, including disk name, path, shelf, bay, serial number, vendor, model, rpm, and carrier serial number.  
**blockNum** (LONGINT): Physical disk block number containing the error.  
**shelf** (STRING): Identifier of the shelf where the disk drive is located.  
**bay** (STRING): Disk bay within the disk shelf where the disk drive is located.  
**vendor** (STRING): Name of the disk drive vendor.  
**model** (STRING): Model string of the disk drive.  
**firmware\_revision** (STRING): Firmware revision number of the disk drive.  
**serialno** (STRING): Serial number of the disk drive.  
**disk\_type** (INT): Type of disk drive.  
**disk\_rpm** (STRING): Rotational speed of the disk drive, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.tetris.lw.id.inconsist

### Severity

NOTICE

## Description

This message occurs during a write operation, when the system detects an inconsistent RAID write signature on a RAID stripe. Data ONTAP® makes appropriate recovery actions. It automatically fails this device safely if the device exceeds the allowed number of inconsistent RAID write signature errors on the disk.

## Corrective Action

(None).

## Syslog Message

Inconsistent RAID write signature detected on RAID group %s%s, stripe #%llu, stripe bitmap %llu, during write operation.

## Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the raid group.

**stripe** (LONGINT): Stripe number.

**more\_stripes** (LONGINT): Bitmap representing all of the stripes that encountered the lost write error. The bit set in this bitmap gives the stripe number relative to the stripes parameter.

## raid.tetris.lw.id.inconsist.dirty

### Severity

NOTICE

### Description

This message occurs during a write operation, when the system detects an inconsistent RAID write signature on a RAID stripe that belongs to a dirty parity region.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**rg** (STRING): Name of the RAID group.

**stripe** (LONGINT): Stripe number.

**more\_stripes** (LONGINT): Bitmap representing all of the stripes that encountered the lost write error. The bit set in this bitmap gives the stripe number relative to the stripes parameter.

## raid.tetris.media.err

### Severity

NOTICE

### Description

This event is issued when tetris i/o detects a medium error

## Corrective Action

(None).

## Syslog Message

Read error on %s%s, block #%llu during stripe write

## Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Formatted information of the disk object that contains the error.  
**blockNum** (LONGINT): The physical block number containing the error.  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.tetris.media.recommend.reassign.err

### Severity

NOTICE

### Description

This message occurs when RAID tetris I/O detects a "recommend reassign" error for a disk block. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

Block recommended for reassignment on %s%s, Block #%llu during stripe write.

### Parameters

**owner** (STRING): Owner of the affected aggregate.  
**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**blockNum** (LONGINT): Physical block number containing the error.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is

located. For non-MetroCluster configurations, site is 'Local'.

## raid.tree events

### raid.tree.remove.activeIOs

#### Severity

NOTICE

#### Description

This message occurs when RAID I/Os are active while removing a disk from the RAID tree. This event is logged in debug builds only.

#### Corrective Action

(None).

#### Syslog Message

Remove disk %s from vol %s: disk has %d active RAID I/Os.

#### Parameters

**disk\_name** (STRING): Name of the disk.  
**tree\_name** (STRING): Name of the RAID tree.  
**active\_ios** (INT): Number of active I/Os to the disk.

## raid.uninitialized events

### raid.uninitialized.parity.vol

#### Severity

EMERGENCY

#### Description

This message occurs periodically to warn the user that the aggregate has no RAID protection and that data loss can occur if a disk fails. Do not fail disks in this aggregate. The state of the aggregate is the result of a special boot mode. If the aggregate requires RAID protection, you should stop writing to it and then restrict it from further writing. Also, avoid creating new aggregates or extending the capacity of existing aggregates. Copy data from existing aggregates to other system and reinitialize the system.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

- WARNING \* %s %s%s was created with no parity; don't fail its data disks!!

#### Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.

# raid.unpartition events

## raid.unpartition.disk

### Severity

NOTICE

### Description

This message indicates a disk was successfully unpartitioned.

### Corrective Action

(None).

### Syslog Message

Disk unpartition successful on %s.

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the disk vendor.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Disk firmware revision number.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.unpartition.disk.fail

### Severity

NOTICE

### Description

This message occurs when an attempt to unpartition a disk fails. The error and error\_info fields indicate the reason why the attempt failed.

### Corrective Action

(None).

### Syslog Message

Disk unpartition failed on %s, error %s, additional error info (%s).

### Parameters

**disk\_info** (STRING): Formatted information of the disk. This includes the disk's containing volume name and other attributes, such as shelf, bay, vendor, model, firmware revision, serial number, and type.

**error** (STRING): String name for the error.

**error\_info** (STRING): Additional error information.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the disk vendor.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Disk firmware revision number.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.unsupported events

### raid.unsupported.bitformat

#### Severity

ERROR

#### Description

This message occurs when configuration management detects that the disk belongs to a 32-bit aggregate, the use of which has been deprecated. The disk is removed from service.

#### Corrective Action

Install a version of Data ONTAP® that supports the 32-bit aggregate, or remove the 32-bit aggregate disks.

#### Syslog Message

The label on %s indicates that the disk belongs to a 32-bit aggregate, the use of which has been deprecated.

#### Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.  
**shelf** (STRING): Identifier of the shelf where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**failure\_reason** (STRING): Text description of the storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.  
**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is Local.

### raid.unsupported.cksumtype

#### Severity

ERROR

## Description

This message occurs when configuration management detects that the disk belongs to a ZCS aggregate, the use of which has been deprecated. The disk is removed from service.

## Corrective Action

Install a version of Data ONTAP® that supports the ZCS aggregate, or remove the ZCS aggregate disks.

## Syslog Message

The label on %s indicates that the disk belongs to a ZCS aggregate, the use of which has been deprecated.

## Parameters

**disk\_info** (STRING): Information about the disk object, including disk name, path, shelf, bay, serial number, vendor, model, RPM, and carrier serial number.

**shelf** (STRING): Identifier of the shelf where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**failure\_reason** (STRING): Text description of the storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is Local.

# raid.verify events

## raid.verify.scrub.aborted

### Severity

NOTICE

### Description

This message occurs when a media scrub is aborted due to a storage layer error. This event is logged in debug builds only.

### Corrective Action

(None).

### Syslog Message

Abort media scrub on %s because of storage layer error: %s.

### Parameters

**disk\_name** (STRING): Name of the disk.

**error** (STRING): Storage layer error.

# raid.vldb events



## raid.vldb.notifyUnsuccessful

### Severity

NOTICE

### Description

This message occurs when RAID fails to notify the M-host that the VLDB(volume location database) is online. The notification is retried automatically when RAID receives VLDB online notification later. To check whether the VLDB process and the mgwd process are up and running, run the "cluster ring show -unitname vldb/mgmt" advanced command.

### Corrective Action

(None).

### Syslog Message

This message occurs when RAID fails to trigger the VLDB online notification in the M-host.

### Parameters

(None).

## raid.vol events

### raid.vol.disk.add.done

### Severity

NOTICE

### Description

This event is issued when a disk is added to a aggregate.

### Corrective Action

(None).

### Syslog Message

Addition of %s to %s %s%s has completed successfully

### Parameters

**disk\_info** (STRING): Formatted information of the disk that has been added to the aggregate.

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate to which the disk is added.

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.vol.disk.grow.done

### Severity

NOTICE

### Description

This message occurs when a drive is grown and is consumed by the owning aggregate by increasing its storage space.

### Corrective Action

(None).

### Syslog Message

%s of %s %s%s has been grown successfully.

### Parameters

**disk\_info** (STRING): Formatted information of the drive that was grown.

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate to which the drive belongs.

**shelf** (STRING): Shelf identifier where the drive is located.

**bay** (STRING): Bay within the shelf where the drive is located.

**vendor** (STRING): Name of the vendor of the drive.

**model** (STRING): Model string of the drive.

**firmware\_revision** (STRING): Firmware revision number of the drive.

**serialno** (STRING): Serial number of the drive.

**disk\_type** (INT): Type of drive.

**disk\_rpm** (STRING): Rotational speed of the drive in RPM.

**carrier** (STRING): Unique ID of the carrier in which the drive is installed.

**site** (STRING): For a MetroCluster® configuration, indicates the site {Local|Remote} where the drive is located. For other configurations, the site is 'Local'.

## raid.vol.failed

### Severity

NOTICE

### Description

This message occurs when the system marks an aggregate as "failed". This results from multiple disks missing or failing in the aggregate. This aggregate is already offline or not serving data.

### Corrective Action

(None).

### Syslog Message

%s %s%s: Failed due to multi-disk error.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## raid.vol.giveback

### Severity

NOTICE

### Description

This message occurs when the system marks an aggregate as "failed" in preparation for sending the aggregate back to its home (owner) node.

### Corrective Action

(None).

### Syslog Message

Aggregate %s%s: Prepared for giveback operation.

### Parameters

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## raid.vol.hybrid

### Severity

INFORMATIONAL

### Description

This message occurs when Solid State Disks (SSDs) are requested to be added for the first time to a hybrid\_enabled aggregate. Addition can be delayed if SSDs need to be zeroed before the addition. If the system reboots or if any of the added SSDs fails before the SSD zeroing is complete, the aggregate will not be converted to a hybrid aggregate.

### Corrective Action

(None).

### Syslog Message

%s%s will become a hybrid %s when the SSDs are added to the aggregate.

### Parameters

**vol\_owner** (STRING): Owner of the hybrid volume.

**vol\_name** (STRING): Name of the hybrid volume.

**vol\_type** (STRING): Volume type.

## raid.vol.inconsist.unmount

### Severity

ERROR

### Description

This message occurs when the system unmounts an aggregate because of file system inconsistency.

### Corrective Action

Run wafllron to clear file system inconsistency. Contact NetApp technical support for assistance.

## Syslog Message

Inconsistent %s %s%s is unmounted.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.vol.mirror.degraded

### Severity

ALERT

### Description

This message occurs when a mirror goes into degraded state as a result of the failure of a plex.

### Corrective Action

To resolve this problem, complete the following steps: 1. Identify the problem volume and failed plex from the system logs. 2. Determine the reason for the plex failure, such as disk drive, disk shelf, or loop failure. 3. Replace any faulty components that you discover. If you need assistance, contact NetApp technical support.

## Syslog Message

%s %s%s is mirrored and one plex has failed. It is no longer protected by mirroring.

## Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## raid.vol.mirror.normal

### Severity

NOTICE

### Description

This message occurs when a mirror that was degraded as a result of a plex failure is no longer degraded. It is also displayed in the following scenarios: 1. When an unmirrored aggregate is mirrored 2. After a resynchronization.

### Corrective Action

(None).

## Syslog Message

%s%s is mirrored and plex is in normal state.

## Parameters

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**aggregate\_uuid** (STRING): Universal Unique Identifier (UUID) of the aggregate.

## **raid.vol.mixed.raid.type**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an aggregate is converted into a mixed RAID type aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

%s%s is now a mixed RAID type aggregate.

### **Parameters**

**vol\_owner** (STRING): Owner of the volume.

**vol\_name** (STRING): Name of the volume.

## **raid.vol.noMirrorSupport**

### **Severity**

ERROR

### **Description**

This event occurs when a mirrored aggregate is detected, but SyncMirror® is not supported.

### **Corrective Action**

Unmirror the aggregate.

### **Syslog Message**

Mirrored %s %s%s is restricted, the system does not support SyncMirror.

### **Parameters**

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

## **raid.vol.online.req.giveback**

### **Severity**

ERROR

### **Description**

This message occurs when the volume could not be brought online because it is being given back to the partner node.

### **Corrective Action**

Bring the volume or aggregate online after giveback is complete.

### **Syslog Message**

Bringing '%s %s' online failed because giveback is in progress.

## Parameters

**vol\_type** (STRING): Whether it is a volume or aggregate that could not be brought online.

**vol** (STRING): Name of the volume or aggregate.

## raid.vol.online.req.nso

### Severity

ERROR

### Description

This message occurs when the volume cannot be brought online because the aggregate containing the volume is part of an ongoing MetroCluster(tm) negotiated switchover source operation.

### Corrective Action

Wait until the switchover is complete, and then bring the volume online by using the "volume online" command. If the "volume online" command fails with a "snapmirror" or "vol move" error, abort the SnapMirror® or volume move operation, and then restart that operation to bring the volume online.

### Syslog Message

Could not bring volume '%s' online because a MetroCluster negotiated switchover source operation is in progress.

## Parameters

**vol** (STRING): Name of the volume.

## raid.vol.online.req.relocate

### Severity

ERROR

### Description

This message occurs when the volume cannot be brought online because the aggregate containing the volume is being relocated to another node.

### Corrective Action

Bring the volume online after a 'storage aggregate relocation' operation is complete.

### Syslog Message

Bringing flexible volume '%s' online failed because 'storage aggregate relocation' operation is in progress.

## Parameters

**vol** (STRING): Name of the volume.

## raid.vol.online.req.sb

### Severity

ERROR

### Description

This message occurs when the volume cannot be brought online because the aggregate containing the volume is part of an ongoing MetroCluster(tm) switchback operation.

## Corrective Action

Wait until the switchback is complete, and then bring the volume online by using the "volume online" command. If the "volume online" command fails with a "snapmirror" or "vol move" error, abort the SnapMirror® or volume move operation, and then restart that operation to bring the volume online.

## Syslog Message

Could not bring volume '%s' online because a MetroCluster switchback operation is in progress.

## Parameters

**vol** (STRING): Name of the volume.

## raid.vol.read.latency.high

### Severity

NOTICE

### Description

This message occurs when the time taken for the completion of a WAFL® disk read operation is more than the specified time limit. This message is logged only on the debug kernel and is used for diagnostic purposes.

### Corrective Action

(None).

### Syslog Message

%s %s%s, vbn #%llu read latency is higher than the specified time limit of %d seconds.

### Parameters

**vol\_type** (STRING): Volume type.

**vol\_owner** (STRING): Owner of the aggregate.

**vol\_name** (STRING): Name of the aggregate.

**read\_vbn** (LONGINT): Volume Block Number of the first block in the read block range that encountered high latency.

**time\_limit** (INT): Time limit, in seconds, that was specified for completion of the read operation.

## raid.vol.reparity.issue

### Severity

NOTICE

### Description

This message occurs when an aggregate has a reparity issue involving NVRAM. Data ONTAP® takes appropriate recovery actions, as described in additional logged events.

### Corrective Action

(None).

### Syslog Message

%s %s%s has %s.

## Parameters

**vol\_type** (STRING): Volume type.  
**owner** (STRING): Owner of the aggregate.  
**vol** (STRING): Name of the aggregate.  
**issue** (STRING): Type of issue.

## raid.vol.replay.nvram

### Severity

INFORMATIONAL

### Description

This event is issued when one or more volume needs nvram replay

### Corrective Action

(None).

### Syslog Message

Performing raid replay on volume(s)

### Parameters

(None).

## raid.vol.root.noMirrorSupport

### Severity

ERROR

### Description

This message occurs when a mirrored root aggregate is detected, but the system does not support SyncMirror®.

### Corrective Action

Unmirror the aggregate.

### Syslog Message

%s aggregate is mirrored, but the system does not support SyncMirror.

### Parameters

**owner** (STRING): Owner of the aggregate.

## raid.vol.rootConflictRestrict

### Severity

NOTICE

### Description

This message occurs when the system forcibly restricts a conflicting root volume, based on a comparison of the name and creation times. This situation could occur if the high-availability (HA) mailbox mirror voting information is lost, following recovery from an HA disaster.



**Corrective Action**

(None).

**Syslog Message**

Volume %s%s is being restricted due to a root volume conflict with %s%s.

**Parameters**

**vol\_owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**rootvol\_owner** (STRING): Owner of the root aggregate.

**rootvol** (STRING): Name of the root volume.

**raid.vol.rootRestrictLessRecent****Severity**

ERROR

**Description**

This message occurs when the system detects multiple root volumes at boot time. Only one root volume should exist per system. A root volume with a less recent creation time is restricted by default. This situation could occur in an upgrade/revert scenario whereby stale/old volumes created on the same system are detected.

**Corrective Action**

If the system default is not the intended root volume, you can switch to the intended root volume.

**Syslog Message**

Volume %s%s is being restricted based on its less recent creation time.

**Parameters**

**rootvol\_owner** (STRING): Owner of the restricted volume.

**rootvol** (STRING): Name of the restricted root volume.

**raid.vol.rootSelectMostRecent****Severity**

ERROR

**Description**

This message occurs when the system detects multiple root volumes at boot time. Only one root volume should exist per system and the system selects the one with the most recent creation time by default. This situation could occur in an upgrade/revert scenario whereby stale/old volumes created on the same system are detected.

**Corrective Action**

If the system default is not the intended root volume, you can switch to the intended root volume.

**Syslog Message**

Volume %s%s is being selected as the default root volume based on its most recent creation time.

## Parameters

**rootvol\_owner** (STRING): Owner of the (default) selected volume

**rootvol** (STRING): Name of the (default) selected root volume

## raid.vol.state.changed

### Severity

NOTICE

### Description

This message occurs the when system detects that the RAID state of an aggregate is changed to either offline or restricted state.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vol\_type** (STRING): Volume type.

**vol\_uuid** (STRING): UUID of the aggregate.

**vol** (STRING): Name of the aggregate.

**state** (STRING): Aggregate state (offline or restricted).

## raid.vol.state.online

### Severity

NOTICE

### Description

This event occurs when the RAID state of an aggregate is moved to the online state.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vol\_type** (STRING): Volume type.

**vol\_uuid** (STRING): UUID of the aggregate.

**vol** (STRING): Name of the aggregate.

## raid.vol.tooBig.allp.reminder

### Severity

ERROR

## Description

This message occurs when the total size of all aggregate plexes is too large.

## Corrective Action

(Call support).

## Syslog Message

The sum of all aggregate plex sizes is %s, exceeding the %s maximum. Contact NetApp technical support.

## Parameters

**volsizes** (STRING): Total size of volumes, excluding the size of all mirrored plexes.

**limit** (STRING): Maximum capacity, excluding total size of mirrored plexes.

## raid.vol.tooBig.allv.reminder

### Severity

ALERT

### Description

This message occurs when the sum of all volume sizes is too large for mirroring.

### Corrective Action

Physically migrate this aggregate to a node with a larger size limit. If you need assistance, contact technical support.

### Syslog Message

The sum of all aggregate sizes is %s, exceeding the %s maximum. Contact NetApp technical support.

### Parameters

**volsizes** (STRING): Total of mirrored volume sizes.

**limit** (STRING): Maximum total of mirrored volume sizes.

## raid.vol.tooBig.ha.reminder

### Severity

ERROR

### Description

This message occurs when the sum of all aggregate sizes is larger than half of the size allowed in a high availability (HA) pair. If the sum of the sizes for this node and its partner exceeds the limit, takeover might fail.

### Corrective Action

Ensure that the sum of all aggregate sizes on this node and its partner does not exceed the limit. Be careful when creating new aggregates and when adding disks to existing volumes and aggregates not to exceed the limit for both nodes in the HA pair. You might need to remove or destroy some aggregates if the limit is exceeded. Disable HA failover using the 'cf disable' command if you cannot bring the sum of all sizes on both nodes below the limit.

### Syslog Message

The sum of all aggregate sizes is %s, exceeding %s (half of the maximum in an HA pair). If the sum of all

aggregate sizes on this node and its partner exceeds the %s maximum, takeover might fail.

### Parameters

**volsizes** (STRING): Total of mirrored aggregate sizes.

**halfpoint** (STRING): Maximum total of mirrored sizes allocated to one node in the HA pair.

**limit** (STRING): Maximum limit of mirrored aggregate sizes.

## raid.vol.tooBig.offline

### Severity

ALERT

### Description

This message occurs when the aggregate cannot be brought online because it is too large.

### Corrective Action

Physically migrate this aggregate to a node with a larger size limit. If you need assistance, contact technical support.

### Syslog Message

%s %s%s cannot be brought online because its size, %s, is larger than the maximum allowed. Contact NetApp technical support.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**volsize** (STRING): Actual volume size.

**limit** (STRING): Maximum volume size.

## raid.vol.tooBig.reminder

### Severity

ERROR

### Description

This event is issued as a reminder that the aggregate size is too large.

### Corrective Action

(Call support).

### Syslog Message

%s size for %s%s %s exceeds limit %s

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**volsize** (STRING): The actual volume size

**limit** (STRING): The maximum volume size

## raid.vol.undestroy.info.missing

### Severity

INFORMATIONAL

### Description

This event is generated when undestroy info raidarea object is missing from one or more disk of a volume or an aggregate.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**disk\_info** (STRING): The formatted information of the disk

**shelf** (STRING): Shelf identifier where the disk is located

**bay** (STRING): Disk bay within the shelf where disk is located

**vendor** (STRING): Name of the vendor of the disk

**model** (STRING): Model string of the disk drive

**firmware\_revision** (STRING): Firmware revision number of the disk

**serialno** (STRING): Serial number of the disk

**disk\_type** (INT): Type of disk drive

**disk\_rpm** (STRING): Rotational speed of disk in RPM

**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

**site** (STRING): For a MetroCluster(tm) configuration, indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

## raid.vol.unprotected.remotesyncmirror

### Severity

ALERT

### Description

This event is issued when we detect a mirrored aggregate with a plex that is not online. In a remotesyncmirror configuration the aggregate won't be available if the plex is not accessible.

### Corrective Action

Bring online the plex that is now offline.

### Syslog Message

%s %s%s is mirrored and one plex is not online. The volume will not be available if a takeover occurs and the online plex is not accessible to the partner node.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

## raid.vol.volinfo.mismatch

### Severity

EMERGENCY

### Description

This message occurs when a mismatch is detected for both copies of the volinfo block during an attempt to mount a SyncMirror aggregate. Data corruption is suspected. The aggregate cannot be mounted. The volinfo block is the superblock in the WAFL® file system. Two copies of the volinfo block are stored in blocks 1 and 2 of the file system. The two copies are the same except when there is an unclean shutdown of the file system, for example, after a panic. In a SyncMirror aggregate, all used blocks are mirrored in two plexes. Volinfo blocks are written so that either both copies of block 1 are the same in two plexes, or both copies of block 2 are the same in two plexes at any point in time. In an unexpected case when that is not true, it is not possible to determine which copy of the volinfo block points to a consistent state of the file system. The message describes a mismatch for only one copy of the volinfo block. A pair of messages is always issued to describe the mismatch for both copies of the volinfo block, because a mismatch for only one copy is not an unusual event.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

%s %s%s: volinfo blocks (VBN %llu) do not match in two plexes. %s: disk %s, magic 0x%x, CP count %d, checksum %u; %s: disk %s, magic 0x%x, CP count %d, checksum %u.

### Parameters

**vol\_type** (STRING): Volume type.

**owner** (STRING): Owner of the aggregate.

**vol** (STRING): Name of the aggregate.

**vbn** (LONGINT): Volume block number of the volinfo block (1 or 2).

**plex1** (STRING): First plex name.

**disk1** (STRING): Name of the disk in the first plex.

**magic1** (INT): Identifier of the block type in the first plex. It should be 0xdab8fbab for the volinfo block.

**CP\_count1** (INT): Consistency point count in the first plex.

**checksum1** (INT): Block checksum in the first plex.

**plex2** (STRING): Second plex name.

**disk2** (STRING): Name of the disk in the second plex.

**magic2** (INT): Identifier of the block type in the second plex. It should be 0xdab8fbab for the volinfo block.

**CP\_count2** (INT): Consistency point count in the second plex.

**checksum2** (INT): Block checksum in the second plex.

## raid.vol.write.latency.high

### Severity

NOTICE

### Description

This message occurs when the time taken for the completion of a WAFL® tetris write operation is more than the specified time limit. This message is logged only on the debug kernel and is used for diagnostic purposes.

## Corrective Action

(None).

## Syslog Message

%s %s%s, vbn #%llu write latency is higher than the specified time limit of %d seconds.

## Parameters

**vol\_type** (STRING): Volume type.

**vol\_owner** (STRING): Owner of the aggregate.

**vol\_name** (STRING): Name of the aggregate.

**write\_vbn** (LONGINT): Volume Block Number of the first block in the write block range that encountered high latency.

**time\_limit** (INT): Time limit, in seconds, that was specified for completion of the write operation.

# raid.warn events

## raid.warn.no.sparecore.disk

### Severity

ERROR

### Description

This message occurs when there are no suitable spare disks left on the node to use for a potential future coredump operation.

### Corrective Action

Assign unowned disks to this node by using the "storage disk assign" command. If there are no unowned disks, physically add more spare disks and assign them to this node.

### Syslog Message

No suitable spare disks available on the node for a potential future coredump operation.

### Parameters

(None).

# raidlm events

## raidlm.cannotmakeprogress events

### raidlm.cannotMakeProgress

#### Severity

ALERT

#### Description

This message occurs when a system does not contain enough spares to correct existing disk failures or misconfigurations. This message is only seen on systems containing disk carriers capable of holding more than one disk. When this message occurs, manual intervention is required because Data ONTAP® cannot automatically correct the existing issues with the system's disk configuration.

#### Corrective Action

To correct this issue on systems that contain carriers with multiple disks, try the following (in the specified order): (1) Assign any unassigned disks that are contained within a multidisk carrier. (2) Add multidisk carriers. (3) If a replacement carrier is available, do the following: a) Fail ALL disks within the carrier. Do not remove a multidisk carrier that contains non-failed disks; doing so results in system degradation and might result in system failure. b) Wait until the disk's amber Attention LEDs stop flashing and are lit continuously. Do not remove a carrier while the shelf OPS panel is alternating between the shelf ID and the Attention Code H1, or while the disks within the carrier are flashing their amber Attention LEDs. c) Replace the carrier. (4) Reassign partner spare disks in the HA configuration. (5) If the problem persists, call NetApp technical support.

#### Syslog Message

Disk %s cannot be evacuated or copied from a multidisk carrier. The system does not contain spares that are suitable for the copy operation.

#### Parameters

**disk** (STRING): Name of the disk in the carrier.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

## raidlm.carrier events

### raidlm.carrier.copy.abort

#### Severity

NOTICE



## Description

This message occurs when Data ONTAP® determines that an ongoing disk copy operation is no longer required and has stopped the disk copy previously started for this disk.

## Corrective Action

(None).

## Syslog Message

Disk %s copy from carrier %s is no longer needed. Operation stopped.

## Parameters

**source\_disk** (STRING): Name of the disk that was being copied.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

## raidlm.carrier.copy.start

### Severity

NOTICE

## Description

This message occurs when the system detects an undesirable RAID layout and Data ONTAP® is attempting to correct the issue by copying a disk to an available spare. The issue is that two disks from the same carrier are part of a single RAID group. If Data ONTAP does not copy one of the disks to an available spare disk, the RAID group would become double degraded if the carrier is removed accidentally. Such an action would result in significant system performance issues and the possibility of exposing the system to data loss.

## Corrective Action

(None).

## Syslog Message

Disk %s will be copied to disk %s to correct the RAID layout.

## Parameters

**source\_disk** (STRING): Name of the disk being copied.  
**target\_disk** (STRING): Name of the copy destination disk.  
**owner** (STRING): RAID group owner.  
**raid\_group\_name** (STRING): RAID group name.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

## raidlm.carrier.evac.abort

### Severity

NOTICE

## Description

This message occurs when Data ONTAP® determines that an ongoing carrier evacuation operation is no longer required and has stopped the disk copy previously started for this disk.

## Corrective Action

(None).

## Syslog Message

Copy operation for disk %s no longer needed. Operation stopped.

## Parameters

**disk** (STRING): Name of the disk in the carrier.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier in which the disk is installed.

## raidlm.carrier.evac.start

### Severity

NOTICE

### Description

This message occurs when a carrier evacuation operation is initiated in a multidisk carrier containing at least one failed disk. Do NOT remove the carrier until the carrier evacuation operation is complete. The carrier evacuation operation is not complete until both of the following occur: the shelf OPS panel stops alternating between the shelf ID and Attention Code H1, and all amber Attention LEDs on the disks stop flashing and are lit continuously.

### Corrective Action

(None).

## Syslog Message

Disk %s is being copied as part of the carrier evacuation process (Shelf Attention Code: H1). Do not remove the carrier until the disk copy is complete.

## Parameters

**disk** (STRING): Name of the disk in the carrier.  
**shelf** (STRING): Disk shelf identifier where the disk is located.  
**bay** (STRING): Disk bay within the shelf where the disk is located.  
**vendor** (STRING): Name of the vendor of the disk.  
**model** (STRING): Model string of the disk.  
**firmware\_revision** (STRING): Firmware revision number of the disk.  
**serialno** (STRING): Serial number of the disk.  
**disk\_type** (INT): Type of disk.  
**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.  
**carrier** (STRING): Unique ID of the carrier being evacuated.

## raidlm.carrier.splitOwner

### Severity

ERROR

## Description

This message occurs when disks within a multidisk carrier belong to different owners. This situation might occasionally occur in takeover or giveback scenarios (partial giveback).

## Corrective Action

All disks within a multidisk carrier should belong to one owner. If the system is in takeover mode, the situation might correct itself when the partner node recovers, or by repeating the giveback. If the system is not in takeover mode, then do the following: (1) Determine which owner should own the disks. (2) If the disk is a spare, assign it to the correct owner. (3) If the disk is not a spare, copy the data to another disk.

## Syslog Message

Disk %s belongs to a different owner than other disks within the carrier.

## Parameters

**disk** (STRING): Name of the disk in the carrier.

**shelf** (STRING): Disk shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where the disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk.

**disk\_rpm** (STRING): Rotational speed of the disk, in RPM.

**carrier** (STRING): Unique ID of the carrier.

# rbac events

## rbac.spuser events

### rbac.spuser.role.notfound

#### Severity

ERROR

#### Description

This message occurs when the role that was earlier associated with the service-processor application user in a previous release was deleted in the current release and is not found during a revert to a previous release.

#### Corrective Action

Re-create the deleted role for the service-processor application user using the command "security login role create".

#### Syslog Message

Role "%s" for service-processor user "%s" was not found on SVM "%s".

#### Parameters

**role\_name** (STRING): Name of the role which was deleted.

**user\_name** (STRING): Name of the user for which the role modification failed.

**vserver\_name** (STRING): Name of the SVM.

# rdb events

## rdb.corruption events

### rdb.corruption.reported

**Severity**

NOTICE

**Description**

This message occurs when the cluster management Data Replication Module determines that the node's local replica set is corrupted. Automatic recovery of this data will be attempted.

**Corrective Action**

(None).

**Syslog Message**

Corruption reported in the local database of Data Replication Module %s, error: %s.

**Parameters**

**unit** (STRING): The unit that is reporting corruption.

**message** (STRING): Error message.

## rdb.env events

### rdb.env.IOFailure

**Severity**

ALERT

**Description**

This message occurs when the cluster management Data Replication Module determines that I/O operations on the node's root volume are failing. This could indicate that the volume is full, not mounted, or any number of other error conditions.

**Corrective Action**

Contact NetApp technical support for assistance with the root volume verification and recovery procedures.

**Syslog Message**

Error encountered in root volume I/O: %s.

**Parameters**

**message** (STRING): Error message.

### rdb.env.processDuplicate

**Severity**

ERROR

## Description

This message occurs when multiple conflicting instances of the management subsystem processes are executing. The situation was safely detected, but it is necessary to understand why it occurred. Normal system startup and operation should ensure that only one of each process type executes on any node. This can occur when an operator inadvertently runs a second instance by hand under the command shell.

## Corrective Action

Verify that all the cluster management processes are running using the diagnostic command 'cluster ring show.' Verify that only one of each is running on the node by using the 'ps' command in the command shell. If this event continues to occur, or if there is any question as to why the extra process was started, then contact NetApp technical support to investigate.

## Syslog Message

Multiple processes executing, detected by advisory locking violation: %s.

## Parameters

**message** (STRING): Error message.

## rdb.env.replicaCorrupt

### Severity

ALERT

## Description

This message occurs when the cluster management Data Replication Module determines that the node's local replica set is corrupted. You must restore the replica set from the other nodes in the cluster, or from backup. If the damage extends beyond the replica set, it might be necessary to recover the root volume.

## Corrective Action

Contact NetApp technical support for assistance with the Data Replication Module and root volume verification and recovery procedures.

## Syslog Message

Corruption in the local database of Data Replication Module: %s.

## Parameters

**message** (STRING): Error message.

## rdb.env.sitelistInvalid

### Severity

ALERT

## Description

This message occurs when the cluster management Data Replication Module detects an inconsistency in the local cluster configuration (the 'sitelist'). This sitelist is rejected, but the situation must be evaluated with the help of NetApp technical support. This usually occurs as a result of a procedural error, for example, a repurposed node without a wipeconfig, a problem with a controller replacement, a conflict with the node's 'setup' parameters, and so on.

### **Corrective Action**

Contact NetApp technical support to investigate the situation.

### **Syslog Message**

The sitelist content does not match the node environment: %s.

### **Parameters**

**message** (STRING): Error message.

## **rdb.ha events**

### **rdb.ha.active**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the management process goes online as the master under two-node 'cluster HA' failover. Any changes to replicated data are synchronized with the peer node after it becomes available.

#### **Corrective Action**

(None).

#### **Syslog Message**

Entering HA-Active: Going online as active master in a high-availability (HA) configuration.

#### **Parameters**

(None).

### **rdb.ha.mboxError**

#### **Severity**

ALERT

#### **Description**

This message occurs when the system detects a problem with the 'cluster HA' on-disk mailboxes, resulting in failover being currently disabled. No data is at risk, but the situation should be addressed. Note: The 'cluster HA' configuration provides a bidirectional failover mechanism for the management processes in a two-node cluster.

#### **Corrective Action**

First ensure that both nodes are booted and operational. Then verify that all the cluster management processes are running on both nodes using the diagnostic command 'cluster ring show'. If there are any problems with this, contact NetApp technical support for guidance. Take the following actions: - Use the 'cluster ha modify -configure false' command to unconfigure high availability (HA). - Use the 'cluster ha modify -configure true' command to reconfigure HA. - Use the 'cluster ha show' diagnostic command to verify that the system is fully configured (both the front-end and on-disk mailboxes should both be 'true').

#### **Syslog Message**

Bidirectional failover under the 'cluster HA' configuration is not currently functional due to problem with the on-disk mailboxes.

**Parameters**

(None).

**rdb.ha.synchronized****Severity**

NOTICE

**Description**

This message occurs when the management process synchronizes its replicated data with its peer node. Bidirectional failover has been reestablished in a 'cluster HA' configuration.

**Corrective Action**

(None).

**Syslog Message**

Entering HA-Normal: Synchronized with peer in a high-availability (HA) configuration.

**Parameters**

(None).

**rdb.internal events****rdb.internal.error****Severity**

ALERT

**Description**

This message occurs when the cluster management Data Replication Module encounters an internal error, possibly causing the process to exit. The process should restart automatically, with full service restored.

**Corrective Action**

Verify that all the cluster management processes are running using the diagnostic command 'cluster ring show.' If not, contact NetApp technical support.

**Syslog Message**

Internal program error encountered: %s.

**Parameters**

**message** (STRING): Error message.

**rdb.node events****rdb.node.starvation****Severity**

ERROR



## Description

This message occurs when the replicated database (RDB) detects CPU starvation on this node, potentially causing it to leave cluster quorum.

## Corrective Action

Reduce the incoming traffic to this node by migrating network interfaces (LIFs) or moving volumes, as necessary.

## Syslog Message

CPU starvation detected in the RDB.

## Parameters

(None).

# rdb.open events

## rdb.open.unit.error

### Severity

EMERGENCY

### Description

This message occurs when the cluster management data replication module encounters an error and fails to open the unit, causing the process to exit. The process should restart automatically with full service restored.

### Corrective Action

Verify that all the units are operational and using the "cluster ring show" diagnostic command. If the problem persists, contact NetApp technical support for assistance.

### Syslog Message

Module "rdb\_module::module\_init": failed to open unit "%s" -and is exiting the process to force a restart.

### Parameters

**unit** (STRING): Unit that failed to open.

# rdb.recovery events

## rdb.recovery.failed

### Severity

EMERGENCY

### Description

This message occurs when the cluster management Data Replication Module could not automatically repair its local replicas.

### Corrective Action

(Call support).

## Syslog Message

Error: %s. Unable to recover the local database of Data Replication Module: %s.

## Parameters

**message** (STRING): Error message.

**unit** (STRING): The unit that failed recovery.

## rdb.recovery.started

### Severity

NOTICE

### Description

This message occurs when the cluster management Data Replication Module starts automatic recovery attempts.

### Corrective Action

(None).

## Syslog Message

Automatic recovery of Data Replication Module %s started.

## Parameters

**unit** (STRING): The unit that is going to start recovery.

## rdb.recovery.succeeded

### Severity

NOTICE

### Description

This message occurs when the cluster management Data Replication Module determines that an automatic recovery was successful.

### Corrective Action

(None).

## Syslog Message

Successful recovery of the local database of Data Replication Module %s.

## Parameters

**unit** (STRING): The unit that finished recovery.

## rdb.sql events

### rdb.sql.upgrade.failed

### Severity

ERROR

## Description

This message occurs when the cluster management Data Replication Module could not upgrade its database from RDB to SQL.

## Corrective Action

Contact Contact NetApp technical support. for assistance.

## Syslog Message

Error: %s. Unable to upgrade the local database from RDB to SQL of Data Replication Module: %s.

## Parameters

**message** (STRING): Error message.

**unit** (STRING): Unit that failed the upgrade.

## rdb.sql.upgrade.succeeded

### Severity

NOTICE

### Description

This message occurs when the cluster management Data Replication Module upgrades its database from RDB to SQL.

### Corrective Action

(None).

### Syslog Message

Successful upgrade of the local database from RDB to SQL of Data Replication Module: %s. Duration: %s sec Tables: %llu Records: %llu

### Parameters

**unit** (STRING): Unit that finished the upgrade.

**duration** (STRING): Duration of the upgrade in seconds.

**tables** (LONGINT): Number of tables processed.

**records** (LONGINT): Number of records processed.

# rdma events

## rdma.rlib events

### rdma.rlib.adv.failed

#### Deprecated

Deprecated as of ONTAP 9.1 because this platform is no longer supported.

#### Severity

ERROR

#### Description

This message occurs when the system fails to start the advertisement service for the high-availability (HA) interconnect.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Failed to start the advertisement service on HA interface %s.

#### Parameters

**if\_name** (STRING): Name.

### rdma.rlib.config.if.failed

#### Deprecated

Deprecated as of ONTAP 9.1 because this platform is no longer supported.

#### Severity

ERROR

#### Description

This message occurs when configuration for the high-availability (HA) interface fails.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Failed to configure the HA interface %s.

#### Parameters

**if\_name** (STRING): Interface name.

### rdma.rlib.get.if.failed

#### Severity

ERROR

**Description**

This message occurs when the system cannot find high availability (HA) capable interface.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

No HA Interconnect capable interface was found.

**Parameters**

(None).

**rdma.rlib.listen.failed****Severity**

ERROR

**Description**

This message occurs when the system fails to start the listener for the HA interface.

**Corrective Action**

Verify that link status is up by running the "metrocluster interconnect adapter show" command. Also check the FC-VI connectivity between DR partners. If the issue persists, contact NetApp technical support.

**Syslog Message**

Failed to start listener for the HA interconnect on interface %s.

**Parameters**

**if\_name** (STRING): Name.

**rdma.rlib.queue.full****Severity**

NOTICE

**Description**

This message occurs when the send or receive queue of a Queue Pair (QP) is full.

**Corrective Action**

(None).

**Syslog Message**

%s queue of QP %s is full.

**Parameters**

**send\_or\_rcv** (STRING): Send or receive queue of a QP.

**qp\_name** (STRING): QP name.

# rdmalib events

## rdmalib.abort events

### rdmalib.abort.nv.reconnect

**Severity**

NOTICE

**Description**

This message occurs when a remote NVLOG operation is aborted due to an error, which is logged here.

**Corrective Action**

(None).

**Syslog Message**

Interconnect channel (%s) encountered error (%s), causing the remote NVLOG to be aborted.

**Parameters**

**channel** (STRING): Interconnect channel.

**err** (STRING): Error that caused the abort.

## rdmalib.duplicateid events

### rdmalib.duplicateid

**Severity**

ALERT

**Description**

This message occurs when the local node system identifier matches the partner node's system identifier.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Interconnect : Partner ID %u is the same as our id.

**Parameters**

**id** (INT): System Id

## rdmalib.multichan events

### rdmalib.multichan.count

**Severity**

NOTICE

## Description

This message occurs when this node's high-availability (HA) partner reports an update on the number of interconnect channels (multichan count) in use.

## Corrective Action

(None).

## Syslog Message

Multichan count "(%u)" in use on HA partner is notified via path "%s".

## Parameters

**count** (INT): Multichan count in use on HA partner.

**notify\_path** (STRING): Path where this event is generated.

# rdmalib.rcm events

## rdmalib.rcm.configMissing

### Severity

NOTICE

### Description

This message occurs when the configuration of the high-availability (HA) interface cannot be completed due to a missing IP address.

### Corrective Action

Verify that the High Availability interconnect card is present on the system and is cabled correctly. If a capable interface is found and the issue persists, contact NetApp technical support.

### Syslog Message

HA configuration is missing, error %d.

### Parameters

**error** (INT): Error code of the missing configuration.

# recover events

## recover.abort events

### recover.abort.ROOLR

**Severity**

NOTICE

**Description**

This event is generated when restore has been aborted and currently ROOLR (Read Only On Line Replica) is running on the Filer. This might be because of the user terminating restore or might be because of bad tape device

**Corrective Action**

(None).

**Syslog Message**

The abort event, %s, is just notified.

**Parameters**

**errmsg** (STRING): The error message is notified

### recover.abort.STSP

**Severity**

NOTICE

**Description**

This event is generated when a recover has been aborted which was started by VERITAS NBU STSP Protocol. This might be because of the user terminating restore or might be because of bad tape device

**Corrective Action**

(None).

**Syslog Message**

The abort event, %s, is just notified.

**Parameters**

**errmsg** (STRING): The error message is notified



# reg events

## reg.cookie events

### reg.cookie.null

#### Severity

ERROR

#### Description

This message occurs when allocation of cookies for traversing all keys under a prefix fails. An EFAULT error indicates that the cause is an empty registry tree. An ENOTSUPP error indicates that the prefix entered might not be supported on a vFiler unit. An ENOMEM error indicates that the system is running low on memory.

#### Corrective Action

For an EFAULT error, check the prefix for spelling errors. For an ENOTSUPP error, try running the command on physical system, rather than a virtual system. For an ENOMEM error, close all applications that are consuming memory, as needed. If the problem persists, contact NetApp technical support

#### Syslog Message

registry: Allocation of cookie for traversing all keys under the prefix %s, failed at function %s. Error number %d: %s.

#### Parameters

**location** (STRING): Location in the registry code where the failure occurred.

**prefix** (STRING): Name of the prefix that was passed as an argument to the function.

**errno** (INT): Internal error message number.

**errmsg** (STRING): Internal error message description.

## reg.errno events

### reg.errno.addAfterReadable

#### Severity

NOTICE

#### Description

This event is issued when Data ONTAP(tm) cannot add an errno-style error to the list of errors. The error must be added before the readable version is generated. No action by the system administrator is needed.

#### Corrective Action

(None).

#### Syslog Message

registry: ignoring attempt to add an error (%d) after errors\_readable was called.

#### Parameters

**errnoval** (INT): The numerical value of the error.

## reg.errno.addInvalid

### Severity

NOTICE

### Description

This event is issued when Data ONTAP(tm) cannot add an errno-style error to the list of errors. The error number must be non-zero or the error string must be non-NULL. No action by the system administrator is needed.

### Corrective Action

(None).

### Syslog Message

registry: in errors\_add\_real, errnoval must be non-zero or str must be non-NULL.

### Parameters

(None).

## reg.file events

### reg.file.createFail

#### Severity

ALERT

#### Description

This message occurs when the system registry file could not be created.

#### Corrective Action

Use the 'df' nodeshell command to check available space on the root volume. If the root volume is full, free some space on the root volume and retry the operation. If the root volume is not full, the system might be running low on memory. Contact NetApp technical support for further assistance.

#### Syslog Message

registry: Cannot create %s file. %s. e.g. Cannot create /etc/registry.local.0 file.

#### Parameters

**file** (STRING): Name of the file.

**errmsg** (STRING): Internal error message explaining why the file could not be created.

### reg.file.linkFail

#### Severity

ERROR

#### Description

Could not create a hard link between system registry files.

### **Corrective Action**

Check whether the source file exists. If the file doesn't exist, contact NetApp technical support.

### **Syslog Message**

registry: Cannot link %s to %s. %s.

### **Parameters**

**sourcefile** (STRING): Name of the source file.

**destfile** (STRING): Name of the destination file.

**errmsg** (STRING): Internal error message explaining why the hard link could not be created.

## **reg.file.openFail**

### **Severity**

ERROR

### **Description**

Could not open a system registry file.

### **Corrective Action**

Check whether the file exists. If it exists, check the file permissions. Make sure it has read and write permissions.

### **Syslog Message**

registry: Cannot open the %s file. %s.

### **Parameters**

**file** (STRING): Name of the file.

**errmsg** (STRING): Internal error message explaining why the file could not be opened.

## **reg.file.renameFail**

### **Severity**

ERROR

### **Description**

This message occurs when the system registry file failed to update. The temporary copy of the registry could not be renamed to the active registry file name.

### **Corrective Action**

Contact NetApp technical support.

### **Syslog Message**

registry: Cannot rename %s to %s. %s. e.g. An error occurred trying to rename /etc /registry.bck.new2 to /etc/registry.bck

### **Parameters**

**sourcefile** (STRING): Name of the source file.

**destfile** (STRING): Name of the destination file.

**errmsg** (STRING): Internal error message explaining why the rename could not be completed.

## reg.file.statsFail

### Severity

ERROR

### Description

Could not get file statistics of the system registry file.

### Corrective Action

Check whether the file exists. If it exists, check the file permissions. Make sure that it has read and write permissions.

### Syslog Message

registry: Cannot get file statistics of %s file. %s.

### Parameters

**file** (STRING): Name of the file.

**errmsg** (STRING): Internal error message explaining why the file statistics could not be obtained.

## reg.file.updateFail

### Severity

ERROR

### Description

This message occurs when the ONTAP® software could not update the system registry file. Updates are made to the registry files to save the changes to the configuration information. Low disk space and low memory can cause this issue.

### Corrective Action

Use the 'df' nodeshell command to check available space on the root volume. If the root volume is full, free some space on the root volume and retry the operation. If the root volume is not full, the system might be running low on memory. Contact NetApp technical support for further assistance.

### Syslog Message

registry: Cannot update %s file.

### Parameters

**file** (STRING): Name of the file.

## reg.file.writeFail

### Severity

ERROR

### Description

Could not write to a system registry file.

### Corrective Action

Use the 'df' command to check available space on the the root volume. If the root volume is full, free some space on the root volume. If the root volume is not full, the system might be running low on memory. Contact NetApp technical support.

## Syslog Message

registry: Cannot write to %s file. %s.

## Parameters

**file** (STRING): Name of the file.

**errmsg** (STRING): Internal error message explaining why data could not be written to the file.

## reg.file.writeMismatch

### Severity

ERROR

### Description

Could not write completely to a system registry file.

### Corrective Action

Use the 'df' command to check available space on the the root volume. If the root volume is full, free some space on the root volume. If the root volume is not full, the system might be running low on memory. Contact NetApp technical support.

## Syslog Message

registry: Cannot write completely to to %s file; expected %d, but wrote %d. %s.

## Parameters

**file** (STRING): Name of the file.

**actualbytes** (INT): Number of bytes that were supposed to be written.

**written** (INT): Number of bytes that were actually written.

**errmsg** (STRING): Internal error message explaining why data could not be written completely to the file.

## reg.key events

### reg.key.invalid

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software cannot access the registry because the key is not valid.

#### Corrective Action

Use a valid key.

#### Syslog Message

registry: Invalid key %s: %s

#### Parameters

**key** (STRING): Registry key that is not valid.

**reason** (STRING): Reason the key is not valid.

**accessID** (STRING): Shows how the registry was accessed.

## reg.key.null

### Severity

ERROR

### Description

This message occurs when the value of key is NULL while the registry set is being implemented on a vFiler(tm) key.

### Corrective Action

(None).

### Syslog Message

vfiler\_registry\_set: Set of key=%s received NULL for value.

### Parameters

**keyName** (STRING): Name of the key.

## reg.options events

### reg.options.cf.change

### Severity

NOTICE

### Description

This message occurs when an option is changed on one cluster node but applies to both the local and partner when the cluster is in takeover mode. The result is that the partner's setting will be overridden during takeover.

### Corrective Action

(None).

### Syslog Message

Option %s changed on one cluster node.

### Parameters

**name** (STRING): Name of the option

### reg.options.cf.mismatch

### Severity

ERROR

### Description

This message occurs during takeover processing when Data ONTAP(tm) discovers that a mismatch exists between the value of an option on the local node and an option on the partner node.

### Corrective Action

(None).

## Syslog Message

Option %s is not the same as on the partner node

### Parameters

**name** (STRING): Name of the option

## reg.options.nullCookie

### Severity

ERROR

### Description

This message occurs when an internal error occurs during the processing of the options portion of the registry.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

registry: Unable to find the option in skiplist datastructure, but the option code is present in ONTAP. Error is %s.

### Parameters

**error** (STRING): Internal error code.

## reg.options.overrideRc

### Severity

ERROR

### Description

This message occurs when Data ONTAP(tm) (or an administrator) sets an option to be a value different than the value in /etc/rc. At the next reboot, the line in /etc/rc will override this setting and a reg.options.overrideRegistry message will be sent.

### Corrective Action

If you want the new option to be persistent, check the /etc/rc file for this key and remove it.

## Syslog Message

Setting option %s to '%s' conflicts with /etc/rc that sets it to '%s'.

### Parameters

**name** (STRING): Name of the option

**newValue** (STRING): New value for the option

**oldValue** (STRING): Existing value in /etc/rc

## reg.transaction events

## reg.transaction.commitFail

### Severity

ERROR

### Description

This message occurs when ONTAP® cannot commit a registry transaction. A system registry transaction is used to update a group of registry keys simultaneously.

### Corrective Action

Check to see whether the system is running low on memory. Contact NetApp technical support.

### Syslog Message

registry: Cannot commit transaction in '%s'. %s

### Parameters

**location** (STRING): Location where the failure occurred.

**errmsg** (STRING): Internal error message explaining why the transaction commit failed.

## reg.transaction.createFail

### Severity

ERROR

### Description

This message occurs when ONTAP® software cannot create a registry transaction. A system registry transaction is used to update a group of registry keys simultaneously.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

registry: Cannot create transaction. Operation: %s.

### Parameters

**operation** (STRING): Registry operation involved in the failure.

## reg.transaction.setkeyFail

### Severity

ALERT

### Description

This message occurs when ONTAP® cannot set a registry key to a value within a registry transaction.

### Corrective Action

Because the transaction failed, modify the registry key manually. Follow the steps in KB article <https://kb.netapp.com/support/index?page=content&actp=LIST&id=S:1015964>.

### Syslog Message

registry: Transaction cannot set "%s" to "%s".



## Parameters

**key** (STRING): Registry key that should have been changed.

**value** (STRING): Value that the key should have been changed to.

# repl events

## repl.archival events

### repl.archival.ts.failure

#### Severity

NOTICE

#### Description

This message occurs when archival tier scan fails for a cloud snapshot. The cloud snapshot will not be scanned for tiering.

#### Corrective Action

(None).

#### Syslog Message

Archival tier scan failure for cloud snapshot "%s" on endpoint %s of volume dsid %llu.

#### Parameters

**snapshot\_id** (STRING): Snapshot UUID.

**ep\_uuid** (STRING): Endpoint UUID.

**volume\_dsid** (LONGINT): Volume dsid.

## repl.autol2l events

### repl.autol2l.unsupported.configuration

#### Severity

INFORMATIONAL

#### Description

This message occurs when the replication destination engine encounters an unsupported configuration during a transfer with storage efficiency.

#### Corrective Action

(None).

#### Syslog Message

Unsupported configuration for relationship '%s'.

#### Parameters

**transferId** (STRING): UUID of the SnapMirror transfer that has the unsupported configuration.

## repl.cg events

## **repl.cg.item.limit.exceeded**

### **Severity**

ERROR

### **Description**

This message occurs when the number of Consistency Group Items involved in a SnapMirror® Granular Sync relationship exceeds its limit in the system.

### **Corrective Action**

Load balance the system to have fewer Consistency Group Items.

### **Syslog Message**

The number of CG Items involved in a SnapMirror Granular Sync relationship exceeded.

### **Parameters**

(None).

## **repl.cg.limit.exceeded**

### **Severity**

ERROR

### **Description**

This message occurs when the number of Consistency Groups(CG) involved in a SnapMirror® Granular Sync relationship exceeds its limit in the system.

### **Corrective Action**

Load balance the system to have fewer Consistency Groups.

### **Syslog Message**

The number of CGs involved in a SnapMirror Granular Sync relationship exceeded.

### **Parameters**

(None).

## **repl.checker events**

### **repl.checker.block.missing**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a SnapMirror® check command detects a missing block on a SnapMirror® destination volume and the volume might be in an inconsistent state that should be fixed as soon as possible.

### **Corrective Action**

Run WAFL check or wafiron as soon as possible to clear this inconsistency and rebuild from the source or resynchronize. If the block is found to be in a Snapshot copy by running 'vol vbncheck' command, deleting

the Snapshot copy should remove the inconsistency. Please contact NetApp technical support for more options.

### Syslog Message

Snapmirror check on volume %s has detected a missing block at volume block number %llu. It is strongly recommended that you run WAFL check or wafiron to clear this inconsistency. Failure to do so might lead to node panic at any time.

### Parameters

**volname** (STRING): Volume name

**vvbn** (LONGINT): Virtual volume block number of the missing block.

## repl.ckpt events

### repl.ckpt.update.generation

#### Severity

INFORMATIONAL

#### Description

This message occurs to record new checkpoint generations subsequent to the first. Transfer will update the current checkpoint if the new checkpoints are smaller.

#### Corrective Action

(None).

### Syslog Message

Volume '%s%s%s' (DSID %llu) received a new checkpoint generation: %llu, current checkpoint generation: %llu.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

**volDsid** (LONGINT): Data set identifier (DSID) of the volume.

**newCkpt** (LONGINT): New checkpoint generation received.

**oldCkpt** (LONGINT): Existing checkpoint generation.

## repl.cloud events

### repl.cloud.incons.buf

#### Severity

ERROR

#### Description

This message occurs when the replication engine encounters an inconsistent block in the cloud.

#### Corrective Action

Contact NetApp technical support.

## Syslog Message

REPL cloud inconsistent: bad block in %s inode (fileid:%s rootcbn:%u) at fbn:%llu level:%u in volume %llu.  
Abort when corruption is encountered:%u.

## Parameters

**file\_space** (STRING): File space.  
**fileid** (STRING): Identifying number for the suspect inode.  
**rootcbn** (LONGINT): Snapshot inofile root.  
**fbn** (LONGINT): File block number.  
**level** (INT): Buffer level.  
**dsid** (LONGINT): Dsid of the volume.  
**abort** (INT): Abort when corrupt block is encountered.

## repl.cloud.incons.gc

### Severity

ERROR

### Description

This message occurs when the replication engine encounters inconsistent map object in the cloud. This does not cause user data corruption but can result in few dangling objects.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

REPL cloud inconsistent: Corrupted cloud data detected in cloud Snapshot copy:%s, cloud Snapshot copy  
UUID:%s, cloud object:%s.

## Parameters

**snapshot\_name** (STRING): Cloud Snapshot copy name.  
**snapshot\_uuid** (STRING): Cloud Snapshot copy Universally Unique ID (UUID).  
**object\_name** (STRING): Cloud object name.

## repl.core.events

## repl.core.nanny.cleanup

### Severity

INFORMATIONAL

### Description

This message occurs when a replication transfer that completed ten or more minutes ago is found and subsequently cleaned up.

### Corrective Action

(None).

## Syslog Message

The replication transfer identified by transfer Id %s was cleaned up since the transfer finished ten or more minutes ago.

## Parameters

**transferid** (STRING): Identifier given for a particular transfer by the replication initiator.

**status** (INT): Progress status of the replication transfer when it was cleaned up.

**failureReason** (INT): Failure reason code if there was a failure. Zero is returned on success.

**failureMsg** (INT): Failure message code if there was a failure. Zero is returned on success.

**failureMsgDetail** (INT): Failure message detail code if there was a failure. Zero is returned on success.

# repl.data events

## repl.data.uncompacted

### Severity

NOTICE

### Description

This message occurs when compacted data is replicated to a volume that does not have data compaction enabled. The data is uncompactd on the destination volume and its original contents are written.

### Corrective Action

(None).

### Syslog Message

Data written to volume %s%s%s will be uncompactd.

## Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

# repl.dest events

## repl.dest.ilabels.notsup

### Severity

ERROR

### Description

This message occurs when a snapshot that contains ilabels is transferred to a destination whose ONTAP version that does not support ilabels.

### Corrective Action

Update the destination node version of ONTAP to 9.9.1 or above.

### Syslog Message

SnapMirror transfer with transfer ID %s will not send existing ilabels because the destination node %s does not support them.

## Parameters

**transferUuid** (STRING): UUID of the SnapMirror® transfer.

**dstNode** (STRING): UUID of the node where the SnapMirror® destination is present.

## repl.dest.multipart.inodes.notsup

### Severity

ERROR

### Description

This message occurs when a Snapshot that contains multi-part inodes is transferred to a destination volume that does not support them.

### Corrective Action

Update the destination's node version of ONTAP that supports multi-part inodes and enable granular data on the destination volume.

### Syslog Message

SnapMirror transfer failed because the transfer Snapshot contains multi-part inodes and the destination volume does not support them.

### Parameters

(None).

## repl.dest.space.guarantee.no.space

### Severity

ERROR

### Description

This message occurs when replication fails to set the space guarantee on the destination volume because there is not enough space on the destination aggregate.

### Corrective Action

Changing the space guarantee type failed because the destination aggregate does not have enough free space. Please increase the amount of free space on the destination aggregate or change the space guarantee type of the source volume to NONE.

### Syslog Message

The volume with DSID %llu, name %s%s, Vserver UUID %s cannot support changing the space guarantee to %s because the aggregate does not have enough free space. Please increase the amount of free space on the destination aggregate by at least %s.

### Parameters

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolName** (STRING): Destination volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume

**srcSpcGuarantee** (STRING): Source volume space guarantee.

**shortBy** (STRING): Amount of extra free space required on the destination aggregate to support changing the guarantee type of the destination volume.

## repl.dest.volume.too.small

## Severity

ERROR

## Description

This message occurs when the transfer Snapshot(tm) copy from the source is unexpectedly larger than the size of the destination volume.

## Corrective Action

Increase the destination volume size to be greater than or equal to the source volume size.

## Syslog Message

The volume with vserver UUID %s, DSID, %llu, name %s, and size %lld bytes, is not large enough to allow replication of a transfer snapshot with size %lld from the source volume.

## Parameters

**dstVolOwnerUUID** (STRING): UUID of the vserver that owns the destination volume.

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolName** (STRING): Destination volume name.

**dstVolSize** (LONGINT): Destination volume size in bytes.

**srcVolSize** (LONGINT): Source transfer snapshot size in bytes.

# repl.dst events

## repl.dst.snaps.check.failed

## Severity

ERROR

## Description

This message occurs when a SnapMirror® logical transfer is not possible due to an incompatible reference Snapshot copy on the destination.

## Corrective Action

Continue with physical transfers until a new Snapshot copy is created on the source volume and the source volume is converted to dual inode space. Logical transfers can be started after this new Snapshot copy is replicated to the destination volume.

## Syslog Message

SnapMirror operation failed because the reference Snapshot copy on the destination volume is either a) from a release prior to Data ONTAP 8.2 b) captured in a nonclustered format c) not upgraded to dual inode space.

## Parameters

**dstVolMsid** (LONGINT): Destination volume Master Data Set ID (MSID).

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolVserverId** (INT): Destination volume Vserver ID.

**dstInstanceReflmageUuid** (STRING): Destination reference instance Snapshot copy UUID.

**dstVersionReflmageUuid** (STRING): Destination reference version Snapshot copy UUID.



# repl.enable events

## repl.enable.seq.split

### Severity

NOTICE

### Description

This message occurs when sequential splitting mode is enabled for a SnapMirror Synchronous relationship. This mode is unsuitable for large file workloads like LUNs, databases, virtualized workloads and video surveillance.

### Corrective Action

(None).

### Syslog Message

Enabling sequential splitting mode for relationship with UUID "%s".

### Parameters

**relationship\_id** (STRING): Relationship UUID.

# repl.exovol events

## repl.exovol.not.quiesced

### Severity

INFORMATIONAL

### Description

This message occurs when an exo-volume scanner detects a loss of volume quiescence during a volume move operation. This typically occurs when there is too much data to transfer within the allowed cutover window.

### Corrective Action

Check the network connectivity between the nodes involved in the volume move. Decrease the amount of client traffic to the volume. If those steps do not reduce the amount of data to be transferred enough to allow the transfer, force a volume move cutover with the "-cutover-action force" flag to the "volume move start" command.

### Syslog Message

(None).

### Parameters

**plugin** (STRING): Exo-volume plug-in name.

**srcVolMsid** (LONGINT): Source volume Master Data Set ID (MSID).

**srcVolDsid** (LONGINT): Source volume Data Set ID (DSID).

**srcVolClusterId** (STRING): Source volume cluster Universally Unique ID (UUID).

**srcVolVserverId** (INT): Source volume Vserver ID.

**bytes\_transferred** (LONGINT): Number of bytes already transferred.

**elapsed\_msecs** (LONGINT): Elapsed time in milliseconds.

# repl.extended events

## repl.extended.time.format.unsupported

### Severity

ERROR

### Description

This message occurs when the replication engine encounters a destination that does not support extended time format.

### Corrective Action

Select a destination that supports extended time format.

### Syslog Message

Extended time format not supported on the destination of relationship %s.

### Parameters

**transferId** (STRING): UUID of the SnapMirror® transfer that has the destination that cannot handle extended time.

# repl.gsr events

## repl.gsr.cache.modify.gs.enfrc.failed

### Severity

ERROR

### Description

This message occurs when SnapMirror faces internal error while going Out-Of-Sync or source volume comes online.

### Corrective Action

Resynchronize SnapMirror sync relation. Please contact NetApp technical support for more options.

### Syslog Message

Unable to modify enforcement for the sync granset relation.

### Parameters

**transferId** (STRING): UUID of the SnapMirror® transfer.  
**srcCgUUID** (STRING): Source CG UUID.  
**Code** (INT): Error code.

## repl.gsr.clone.replication

### Severity

NOTICE

### Description

This message occurs when replication of a clone request fails or cannot be serviced. The clone request

completed successfully on the sync mirror primary, but was not completed on the sync mirror secondary. The clone will not be available on the sync mirror secondary.

### Corrective Action

For further information about correcting the problem, see the Knowledgebase article #.

### Syslog Message

Failed to replicate clone of /vol/%s/%s to sync mirror secondary.

### Parameters

**volname** (STRING): Volume name.

**filename** (STRING): Source filename.

**reason** (STRING): Descriptive reason for the notice.

## repl.gsr.snap.replication.failed

### Severity

INFORMATIONAL

### Description

This message occurs when a Snapshot® replication request fails or cannot be served.

### Corrective Action

(None).

### Syslog Message

Failed to replicate snapshot operation to the destination.

### Parameters

**transferId** (STRING): UUID of the SnapMirror® transfer.

**srcVolumeDsid** (LONGINT): Source volume data set ID (DSID).

**srcCgUUID** (STRING): Source CG UUID.

**snapshotName** (STRING): Name of the Snapshot copy to be replicated.

**failureReason** (STRING): Reason for the failure.

**replStatus** (INT): Replication Status.

**replFailureMsg** (INT): Replication failure message.

**replFailureMsgDetail** (INT): Replication failure message detail.

## repl.holdblock events

### repl.holdblock.arrived

### Severity

INFORMATIONAL

### Description

This message occurs during an incremental transfer when the replica volume did a snapshot restore before the transfer. It records the VVBN of the block which is supposed to be held in memory until the end of transfer as it arrives from the source.

## Corrective Action

(None).

## Syslog Message

Held block for volume '%s%s%s' (DSID %llu) with VVBN %llu has arrived.

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

**volDsid** (LONGINT): Data set identifier (DSID) of the volume.

**vvbn** (LONGINT): VVBN of the held block.

## repl.holdblock.begin

### Severity

INFORMATIONAL

### Description

This message occurs at the beginning of an incremental transfer if the replica volume did a snapshot restore. It records the VVBN of the block that needs to be held in memory until the end of transfer.

### Corrective Action

(None).

### Syslog Message

Holding block for volume '%s%s%s' (DSID %llu) with VVBN %llu because Revert-to snapshot is set to %d.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

**volDsid** (LONGINT): Data set identifier (DSID) of the volume.

**vvbn** (LONGINT): VVBN of the block to be held.

**snapld** (INT): Revert-to snapshot Id.

## repl.holdblock.end

### Severity

INFORMATIONAL

### Description

This message occurs at the end of an incremental transfer when the replica volume did a snapshot restore. The required VVBN arrived from the source and has been applied to WAFL.

### Corrective Action

(None).

### Syslog Message

Applied held block for volume '%s%s%s' (DSID %llu) with VVBN %llu to WAFL.

## Parameters

**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier for the volume.  
**volDsid** (LONGINT): Data set identifier (DSID) of the volume.  
**vvbn** (LONGINT): VVBN of the held block applied to WAFL.

# repl.hole events

## repl.hole.raid.error

### Severity

ALERT

### Description

This message occurs when a SnapMirror® transfer encounters an unrecoverable RAID error.

### Corrective Action

Contact NetApp technical support for further assistance.

### Syslog Message

SnapMirror transfer with transfer ID %s encountered an unrecoverable RAID error for %s with ID %d in destination volume %d.

## Parameters

**transferId** (STRING): UUID of the SnapMirror transfer that encountered a RAID error.  
**fileSpace** (STRING): File space (metafile or user file).  
**fileId** (INT): File ID of the inode involved if it is a user file space, and metafile name if it is a metafile space.  
**dstVolDsid** (LONGINT): Data set identifier (DSID) of the destination volume.

# repl.jumpahead events

## repl.jumpahead.failed

### Severity

ERROR

### Description

This message occurs when the system cannot bring a replica volume online after applying the latest update to the replica because of possible inconsistencies detected in the update. The system rolls back the volume to the file system that existed prior to applying the update and brings the volume back online.

### Corrective Action

Run wafiron as soon as possible on the aggregate where the replica volume resides as well as on the aggregate where the source volume of the replica resides. If the problem persists, contact NetApp technical support for further assistance.

### Syslog Message

Rolling back after jump-ahead failure for volume '%s%s%s' (DSID %llu) with error %d.

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume

**volDsid** (LONGINT): Data Set Identifier (DSID) of the volume

**code** (INT): Error code returned while bringing online the volume with the latest update

# repl.large events

## repl.large.vol.600tb.limits.unsup

### Severity

ERROR

### Description

This message occurs when the replication engine encounters a destination that cannot handle large 600TB volume limits.

### Corrective Action

Take one of the following actions: - Select a destination with version ONTAP 9.14.0 or later. - On the source volume, reduce the volume size "volume modify -vserver vserver\_name> -volume vol\_name -max-size 300TB".

### Syslog Message

Large volume limits not supported on the destination node %s for destination volume with MSID %llu of relationship %s.

## Parameters

**dstNode** (STRING): UUID of the node where the SnapMirror® destination is present.

**dstVolMsid** (LONGINT): MSID of the SnapMirror® destination volume.

**transferId** (STRING): UUID of the failed SnapMirror® transfer.

## repl.large.vol.limits.unsup

### Severity

ERROR

### Description

This message occurs when the replication engine encounters a destination that cannot handle large volume limits.

### Corrective Action

Take one of the following actions: - Select a destination with version ONTAP 9.12.1 or later. - On the source volume, reduce the volume size "volume modify -vserver vserver\_name> -volume vol\_name -large-volume -file-limits disable". Now disable large volume limits by using the "volume modify -vserver vserver\_name -volume vol\_name -large-volume-file-limits disable" command.

### Syslog Message

Large volume limits not supported on the destination node %s for destination volume with MSID %llu of relationship %s.

## Parameters

**dstNode** (STRING): UUID of the node where the SnapMirror® destination is present.

**dstVolMsid** (LONGINT): MSID of the SnapMirror® destination volume.

**transferId** (STRING): UUID of the failed SnapMirror® transfer.

## repl.large.vol.unsup.maxvbn

### Severity

ERROR

### Description

This message occurs when the replication engine encounters a destination that cannot handle large volumes.

### Corrective Action

Select a destination that supports a volume size greater than 100 TB.

### Syslog Message

Large volumes not supported on the destination node %s for destination volume with MSID %llu of relationship %s.

## Parameters

**dstNode** (STRING): UUID of the node where the SnapMirror® destination is present.

**dstVolMsid** (LONGINT): MSID of the SnapMirror® destination volume.

**transferId** (STRING): UUID of the SnapMirror® transfer that has the destination that cannot handle large volumes.

## repl.large.vol.unsup.nomsiz

### Severity

ERROR

### Description

This message occurs when the replication engine encounters a destination that cannot handle large volumes.

### Corrective Action

Take one of the following actions: - On the source volume, disable autosize if it is enabled by using the "volume modify vol\_name -autosize-mode" command. Then reduce the size of the volume by using the "volume modify vol\_name -size" command. - On the source volume, set the "max autosize" parameter to below 100 TB by using the "volume modify vol\_name -max-autosize" command. - On the source volume, reduce the size of the volume to below 100 TB by using the "volume modify vol\_name -size" command. - Select a destination that supports a volume size greater than 100 TB.

### Syslog Message

Large volume not supported on the destination node %s for destination volume with MSID %llu of relationship %s.

## Parameters

**dstNode** (STRING): UUID of the node where the SnapMirror® destination is present.

**dstVolMsid** (LONGINT): MSID of the SnapMirror® destination volume.

**transferId** (STRING): UUID of the failed SnapMirror® transfer.

# repl.logical events

## repl.logical.dw.rebuild

### Severity

INFORMATIONAL

### Description

This message occurs when a data warehouse is rebuilt as part of a SnapMirror® restore or resync operation.

### Corrective Action

(None).

### Syslog Message

A data warehouse rebuild was performed on volume with name %s, MSID %llu, DSID %llu, and Vserver ID %u.

### Parameters

**volName** (STRING): Volume name.  
**volMsid** (LONGINT): Volume master data set ID (MSID).  
**volDsid** (LONGINT): Volume data set ID (DSID).  
**volVserverId** (INT): Volume Vserver ID.

## repl.logical.lreTolrse

### Severity

INFORMATIONAL

### Description

This message occurs when previously enabled compression is disabled on the destination volume of a SnapMirror® relationship with the type XDP, and when it is followed by the "snapmirror update" command with the "enable-storage-efficiency" parameter set to "true". All future transfers to this destination volume from the current source will preserve storage efficiency as long as compression is disabled on it.

### Corrective Action

(None).

### Syslog Message

All future transfers from the current source to destination volume %s%s with MSID %llu, DSID %llu, from the current source will preserve storage efficiency as long as compression is disabled on the destination volume.

### Parameters

**dstVolName** (STRING): Destination volume name.  
**dstVolVserverUuid** (STRING): Destination volume Vserver ID.  
**dstVolMsid** (LONGINT): Destination volume Master data set ID (MSID).  
**dstVolDsid** (LONGINT): Destination volume data set ID (DSID).



## repl.logical.IrseTolre

### Severity

INFORMATIONAL

### Description

This message occurs when compression is enabled on the destination volume of a SnapMirror® relationship that has type XDP. All future transfers to this destination volume from the current source will not preserve storage efficiency.

### Corrective Action

(None).

### Syslog Message

All future transfers to destination volume with MSID %llu, DSID %llu, and Vserver ID %u, from the current source will not preserve storage efficiency.

### Parameters

**dstVolMsid** (LONGINT): Destination volume Master Data Set ID (MSID).

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolVserverId** (INT): Destination volume Vserver ID.

## repl.logical.rawDirToLogicalDir

### Severity

INFORMATIONAL

### Description

This message occurs when a SnapMirror® transfer uses logical mode after a previous transfer to the same destination used raw directory mode. This most frequently occurs when the source node of the transfer was upgraded to a newer version of Data ONTAP® than the destination node.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**dstVolMsid** (LONGINT): Destination volume Master Data Set ID (MSID).

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolClusterId** (STRING): Destination volume cluster Universally Unique ID (UUID).

**dstVolVserverId** (INT): Destination volume Vserver ID.

**srcInstancelmageUuid** (STRING): Source transfer instance Snapshot(tm) copy UUID.

**srcVersionImageUuid** (STRING): Source transfer version Snapshot copy UUID.

**rawDirsSupported** (INT): Indicates whether the destination believes raw directory transfer is supported.

**rawDirsRequired** (INT): Indicates whether the destination believes raw directory transfer is required.

## repl.out events

## repl.out.of.band.vbn

### Severity

INFORMATIONAL

### Description

This message occurs when replication source is asked to send certain blocks out of band because the replica volume did a snapshot restore.

### Corrective Action

(None).

### Syslog Message

Asked volume '%s%s%s' (DSID %llu) to send out-of-band VVBN %llu; current cursor: %llu; need-many-VBN value: %d.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

**volDsid** (LONGINT): Data set identifier (DSID) of the volume.

**oobVVBN** (LONGINT): Out-of-band Virtual Volume Block Number (VVBN).

**cursor** (LONGINT): Current scan cursor.

**many** (INT): Whether multiple out-of-band VVBNs were asked for.

## repl.out.of.memory

### Severity

NOTICE

### Description

This message occurs when an out-of-memory error is encountered during a SnapMirror® replication operation.

### Corrective Action

(None).

### Syslog Message

SnapMirror replication transfer encountered an out-of-memory error.

### Parameters

**memoryAllocation** (STRING): Type of memory allocation.

**retry\_time\_msecs** (LONGINT): Retry time in milliseconds.

**filename** (STRING): Name of the file where the event occurred.

**functionName** (STRING): Name of the function for which the error occurred.

**lineNumber** (INT): Line number at which error the occurred.

**size** (LONGINT): Size of the memory allocation.

## repl.physdiff events

## repl.physdiff.invalid.hole

### Severity

EMERGENCY

### Description

This message occurs when a block to replicate was found to be inconsistent in the source volume. Replication has been aborted to avoid propagating the inconsistency to the destination node.

### Corrective Action

Run wafliiron on the aggregate containing the source volume of the replication.

### Syslog Message

Replication of volume '%s%s%s' (DSID %llu) encountered an inconsistent block (VVBN %llu).

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume.

**volDsid** (LONGINT): Data set identifier (DSID) of the source volume.

**vvbn** (LONGINT): Virtual Volume Block Number of the affected block.

## repl.s2c events

### repl.s2c.checksum.mismatch

#### Severity

ERROR

#### Description

This message occurs when the checksum of an object's contents does not match with the expected checksum.

#### Corrective Action

Please contact NetApp technical support for more options.

#### Syslog Message

Checksum mismatch in cloud object contents. Endpoint UUID = %s. Relationship UUID = %s. Object Name = %s. Slot ID = %d.

#### Parameters

**epUuid** (STRING): Endpoint UUID.

**relnUuid** (STRING): Relationship UUID.

**objName** (STRING): Cloud object name.

**slotId** (INT): Slot ID.

### repl.s2c.rel.state.inconsistent

#### Severity

EMERGENCY

## Description

This message occurs when the replication source engine encounters a disk error in the relationship state metafile. Replication transfer will abort. Current relationship cannot be resumed.

## Corrective Action

Current relationship cannot be resumed. Please create a new relationship.

## Syslog Message

relationship state metafile inconsistent in volume '%s%s%s' due to disk errors. Aborting transfer.

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## repl.s2c.restore.unsupported.config

### Severity

ERROR

### Description

This message occurs when we encounter unsupported configuration during SnapMirror® restore.

### Corrective Action

Cannot continue with restore.

### Syslog Message

Cannot restore to volume '%s%s%s' because %s.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Specific reason for failure.

## repl.s2c.restore.unsupported.snapshot

### Severity

ERROR

### Description

This message occurs when we encounter unsupported source snapshot copy during SnapMirror® restore.

### Corrective Action

Cannot use the source snapshot copy for SnapMirror® restore.

### Syslog Message

Cannot restore from source snapshot copy to volume '%s%s%s' because %s.

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Specific reason for failure.

## repl.s2c.rst.archv.data.fnd

### Severity

NOTICE

### Description

This message occurs when SnapMirror® restore encounters data present in the cloud archive tier.

SnapMirror restore cannot proceed until archival tier data processing has completed.

### Corrective Action

Use the "snapmirror show" command to view the status of the processing of the archival tier of the cloud.

When the relationship status becomes "Idle", then reissue the SnapMirror restore request.

### Syslog Message

SnapMirror restore with transfer UUID "%s", cloud source endpoint UUID "%s", destination volume DSID "%llu", encountered data in the archival tier.

## Parameters

**transferId** (STRING): SnapMirror transfer universal unique identifier (UUID).

**srcEpUuid** (STRING): Cloud source endpoint UUID.

**dstVolDsid** (LONGINT): Restore destination volume data set ID (DSID).

## repl.s2c.src.inconsistent

### Severity

EMERGENCY

### Description

This message occurs when the replication source engine encounters a disk error in the file system metadata. Replication transfer will abort.

### Corrective Action

Run wafiron to fix the active file system. To continue replication transfer, run snapmirror abort -h command on the relationship, delete the current transfer snapshot copy and resume replication transfer from a new clean snapshot copy.

### Syslog Message

WAFI metadata inconsistent in volume '%s%s%s' due to disk errors. Aborting transfer.

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## repl.s2c.unsupported.config

### Severity

ERROR

### Description

This message occurs when we encounter unsupported configuration during SnapMirror® backup.

### Corrective Action

Cannot continue with backup.

### Syslog Message

Cannot backup from volume '%s%s%s' because %s.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Specific reason for failure.

## repl.s2c.vmap.inconsistent

### Severity

INFORMATIONAL

### Description

This message occurs when the replication source engine encounters a disk error in the replication specific metadata. Replication transfer will abort.

### Corrective Action

(None).

### Syslog Message

Replication specific metadata inconsistency in volume '%s%s%s'. Aborting transfer. Transfer will attempt baseline to recover from this condition.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## repl.send events

## repl.send.mgr.forcedLogical

### Severity

INFORMATIONAL

### Description

This message occurs when a SnapMirror® transfer uses logical mode without even attempting raw directory negotiation.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**srcVolMsid** (LONGINT): Source volume Master Data Set ID (MSID).

**srcVolDsid** (LONGINT): Source volume Data Set ID (DSID).

**srcVolClusterId** (STRING): Source volume cluster Universally Unique ID (UUID).

**srcVolVserverId** (INT): Source volume Vserver ID.

**srcInstanceImageUuid** (STRING): Source transfer instance Snapshot(tm) copy UUID.

**srcVersionImageUuid** (STRING): Source transfer version Snapshot copy UUID.

**srcInstanceReflmageUuid** (STRING): Source reference instance Snapshot copy UUID.

**srcVersionReflmageUuid** (STRING): Source reference version Snapshot copy UUID.

# repl.seq events

## repl.seq.split.enabled.SAN

### Severity

ALERT

### Description

This message occurs when the SnapMirror sequential split is enabled on a SAN operating workload. This is meant to be enabled for NAS workloads only. SAN operating workload would experience a higher latency and is not recommended to be used with sequential splitting mode.

### Corrective Action

Switch back to parallel split mode by running the command "snapmirror policy modify -policy (policy name) -are-data-ops-sequentially-split false".

### Syslog Message

Sequential splitting mode is enabled in SAN operating workload for relationship with UUID "%s".

### Parameters

**relationship\_id** (STRING): SnapMirror relationship universal unique identifier (UUID).

# repl.ser events

## repl.ser.ops.pardir.breach

### Severity

NOTICE

### Description

This message occurs when number of operations serialized on parent directory exceed an internal limit. This does not harm the system/cluster in any way.

## Corrective Action

(None).

## Syslog Message

Number of operations serialized on parent directory exceed the internal limit for relationship with UUID %s.

## Parameters

**relationship\_id** (STRING): SnapMirror relationship universal unique identifier (UUID).

# repl.space events

## repl.space.guarantee.transfer.failed

### Severity

ERROR

### Description

This message occurs when the attempt to change space guarantee type fails on the destination volume.

### Corrective Action

Change the space guarantee type of the source volume to be the same as the space guarantee type of the destination volume.

### Syslog Message

The volume with DSID %llu, name %s%s, Vserver UUID %s cannot support changing the space guarantee to %s on the destination aggregate, error %d.

### Parameters

**dstVolDsid** (LONGINT): Destination volume Data Set ID (DSID).

**dstVolName** (STRING): Destination volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume

**srcSpcGuarantee** (STRING): Source volume space guarantee.

**code** (INT): Error code returned when setting the space guarantee of the destination volume.

# repl.src events

## repl.src.snaps.check.failed

### Severity

ERROR

### Description

This message occurs when SnapMirror® logical transfer is not possible due to incompatible Snapshot copies.

### Corrective Action

Create new Snapshot copy on the source volume and use that for Logical transfer.



## Syslog Message

SnapMirror operation failed because the reference and/or the transfer Snapshot copies on the source volume are either a) from a release prior to Data ONTAP 8.2 b) captured in a nonclustered format c) not upgraded to dual inode space. These Snapshot copies cannot be used for logical transfer.

## Parameters

**srcVolMsid** (LONGINT): Source volume Master Data Set ID (MSID).

**srcVolDsid** (LONGINT): Source volume Data Set ID (DSID).

**srcVolVserverId** (INT): Source volume Vserver ID.

**srcInstanceImageUuid** (STRING): Source transfer instance Snapshot copy UUID.

**srcVersionImageUuid** (STRING): Source transfer version Snapshot copy UUID.

**srcInstanceReflmageUuid** (STRING): Source reference instance Snapshot copy UUID.

**srcVersionReflmageUuid** (STRING): Source reference version Snapshot copy UUID.

# repl.switch events

## repl.switch.to.seq.split

### Severity

NOTICE

### Description

This message occurs when the average Op Latency increases beyond the threshold value in parallel splitting mode.

### Corrective Action

Switch to sequential splitting mode by setting the SnapMirror policy option "-are-data-ops-sequentially-split".

## Syslog Message

Operation Latency has exceeded the threshold value "%llu" in the parallel split mode for sync relationship with UUID "%s".

## Parameters

**average\_overall\_latency** (LONGINT): Average overall latency.

**relationship\_id** (STRING): Relationship UUID.

# repl.unable events

## repl.unable.to.create.css

### Severity

NOTICE

### Description

This message occurs when the workflow is unable to create a common Snapshot copy due to timeout. This operation is retried automatically.

### Corrective Action

(None).

## Syslog Message

Unable to create CSS "%s" for relationship with UUID %s.

## Parameters

**snapshot\_name** (STRING): Common Snapshot copy name

**relationship\_id** (STRING): SnapMirror relationship universal unique identifier (UUID).

# repl.xfer events

## repl.xfer.term.timed.out

### Severity

ERROR

### Description

This message occurs when a replication transfer component (sender or receiver) has taken more than an hour to terminate. Any future manual or scheduled replication transfer for this relationship cannot begin until this condition is corrected.

### Corrective Action

Please generate a core file on the node, submit it to NetApp technical support and reboot the node.

## Syslog Message

A component of a replication transfer with destination volume DSID '%llu' or SnapMirror® tape ndmp session '%s' is taking a long time to terminate since it is busy performing some operation.

## Parameters

**dstVoldsid** (LONGINT): Data set identifier (DSID) of the destination volume.

**smTapeNDMPSession** (STRING): NDMP session identifier for SnapMirror® tape transfers.

# replaycache events

## replaycache.arcbktallocerr events

### ReplayCache.ArcBktAllocErr

#### Severity

ERROR

#### Description

This message occurs when the replay-cache bucket (which is used to hold the responses of non-idempotent operations, to prevent the execution of duplicate I/O requests) that is associated with a bin (through a client connection) has fully utilized the maximum allowed memory and old bins holding the bucket is not able to release requests due to the configured expiration time. I/O operations are suspended until the bucket has available memory.

#### Corrective Action

Modify the replay cache bin expiration time using the "spinnp-replay-cache attributes modify -arc-initial -bucket-count 16 -arc-bin-expiration 30" command (privilege: diag). Specify a smaller value for the bin expiration time and initial bucket count, so that the bin eviction will happen sooner and new bin creation will require less memory in the bucket, which might help avoid the issue.

#### Syslog Message

A replay-cache bucket memory exhaust occurred due to the slow eviction of old bins and the high number of I/O operations being received by the server.

#### Parameters

(None).

## replaycache.nfsv41bindelete events

### ReplayCache.Nfsv41BinDelete

#### Severity

NOTICE

#### Description

This message occurs when the replay-cache bin associated with an NFSv4.1 session is destroyed due to a policy mismatch. In Data ONTAP® releases prior to 9.0, replay-cache bins for NFSv4.1 session were created with a certain policy. The policy associated with these bins has changed and the data in older bins cannot be migrated. Older bins are destroyed and new bins are created. The process of bin policy conversions for NFSv4.1, occurs when the effective cluster version (ECV) changes due to an upgrade or revert operation. Connection Session Manager (CSM) sessions are used to communicate file operations between nodes in the cluster. During this period, if CSM sessions experience timeouts or other errors, NFSv4.1 clients might receive unexpected results for a few of their outstanding requests. All new NFSv4.1 client requests are unaffected.

#### Corrective Action

(None).

**Syslog Message**

The replay-cache bin associated with a NFSv4.1 session was deleted and recreated due to a policy mismatch. Current NFSv4.1 session might temporarily receive unexpected results.

**Parameters**

(None).

# replication events

## replication.cluster events

### replication.cluster.vol

#### Severity

ERROR

#### Description

This message occurs because the transfer of clustered volumes is not supported in nonclustered environments.

#### Corrective Action

Use clustered Data ONTAP to transfer cluster volumes.

#### Syslog Message

Canceled transfer of %s volume '%s', %s transfer of cluster volume is not supported in nonclustered environments.

#### Parameters

**type** (STRING): Destination or source.

**volume** (STRING): Replica volume.

**transfertype** (STRING): Type of transfer.

## replication.dst events

### replication.dst.attrsNotSaved

#### Severity

ERROR

#### Description

This message is issued at the end of a Qtree SnapMirror® or SnapVault® transfer or checker/fixer operation when the replication destination detects that the replication source has sent file attributes that the replication destination does not support, and the replication destination has not saved those attributes. This may happen when the replication source is running a newer version of Data ONTAP® which supports file attributes that the replication destination does not support.

#### Corrective Action

To store all of the attributes, upgrade your replication destination filer to a version of Data ONTAP which is at least as recent as the source. Alternatively, this message can be suppressed for all future transfers or checker/fixer operations for all Qtree SnapMirror and SnapVault relationships by setting the hidden replication option, `suppress_attrs_not_saved_warnings`.

#### Syslog Message

The replication destination was unable to replicate all file attributes on %d file(s) or director(y/ies) in %s because it does not support one or more of the attributes sent by the source.

## Parameters

**attrs\_not\_saved** (INT): The number of files for which the replication destination could not replicate all file attributes.  
**dst\_path** (STRING): Destination qtree path.

## replication.dst.debugLog

### Severity

ERROR

### Description

This message occurs when a SnapMirror® or SnapVault® destination cannot log the replication protocol header for the data received.

### Corrective Action

The reasons for the error cited in the error message might indicate ways to address the problem.

### Syslog Message

Logical Replication header log: %s %s:%s.

## Parameters

**reason** (STRING): Reason for the failure.  
**file\_name** (STRING): Logging file for which the error was encountered.  
**errstr** (STRING): Error encountered during logging operation.

## replication.dst.err

### Severity

ERROR

### Description

This message occurs when the system detects an error on a SnapMirror® or SnapVault® destination volume or qtree.

### Corrective Action

(None).

### Syslog Message

%s: destination transfer from %s:%s to %s : %s.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).  
**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source volume or qtree.  
**dstpath** (STRING): Replication destination volume or qtree.  
**error** (STRING): Specific error message.

## replication.dst.inodeMissing

**Severity**

ERROR

**Description**

This message occurs when a file for which the replication source sent data is not present on the replication destination.

**Corrective Action**

Check for WAFL® inconsistencies on the replication source.

**Syslog Message**

The replication destination could not find the file corresponding to %d sent from the replication source. Check for WAFL inconsistencies on the replication source %s:%s.

**Parameters**

**inode** (INT): Inode number of the source file for which the corresponding file is not present.

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**replication.dst.langMismatch****Severity**

ERROR

**Description**

This message occurs when the user continues with the updates of the replication relationship after changing the language setting of the volume in which qtree replicas are configured or that contain replica primary.

**Corrective Action**

Create a new baseline for the relationship.

**Syslog Message**

Volume language was changed before this transfer or it was changed before some previous update transfer for the relationship of qtree %s. This might create some problems with further replication transfers of the relationship.

**Parameters**

**qtpath** (STRING): Destination qtree path.

**replication.dst.lostACL****Severity**

ERROR

**Description**

This message occurs when a replication update discovers that an inode that is supposed to have an ACL does not have one, causing the replication to fail.

**Corrective Action**

If the replication destination is a SnapMirror® destination, run the 'snapmirror quiesce' command to quiesce

the replication relation, then run the 'snapmirror check -F' command under the guidance of NetApp technical support. If the destination is a SnapVault® destination, run the 'snapvault check -F' command under the guidance of NetApp technical support.

### Syslog Message

Replication transfer from %s:%s to %s:%s: the destination could not find ACL %d for source inode %d.

### Parameters

**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source path.  
**dstfiler** (STRING): Replication destination system.  
**dstpath** (STRING): Replication destination path.  
**acl\_inode** (INT): Inode number of source ACL.  
**file\_inode** (INT): Inode number of the replication source file.

## replication.dst.noCommonSnap

### Severity

ERROR

### Description

This message occurs when a SnapVault® or SnapMirror® relationship cannot be resynchronized because there is no usable Snapshot(tm) copy that the source and destination have in common. In such situations, the 'snapvault start -r' and 'snapmirror resync' commands cannot be successfully completed.

### Corrective Action

The relationship cannot be resynchronized so it must be reinitialized. For SnapMirror, use the 'snapmirror initialize' command. For SnapVault, use the 'snapvault start' command. Either run the baseline transfer to a new qtree, or remove the existing qtree before starting the transfer. Refer to the Data Protection Guide for details.

### Syslog Message

%s: destination transfer from %s:%s to %s:%s : no common Snapshot copy to use as the base for resynchronization.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).  
**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source path.  
**dstfiler** (STRING): Replication destination system.  
**dstpath** (STRING): Replication destination path.  
**rplerror** (STRING): Replication error message.  
**waferror** (STRING): WAFL® error

## replication.dst.noMemory

### Severity

ALERT

### Description

This message occurs when there is insufficient memory on the SnapMirror® or SnapVault® destination for the transfer to continue.



## Corrective Action

(None).

## Syslog Message

System is low on memory; aborting %s destination transfer.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

## replication.dst.noSrcSnap

### Severity

ERROR

### Description

This message occurs when a SnapVault® or SnapMirror® destination discovers that the relationship's base Snapshot(tm) copy no longer exists on the source. In such situations, 'snapvault update' and 'snapmirror update' commands cannot execute successfully.

### Corrective Action

For SnapMirror, issue the command 'snapmirror resync' to make SnapMirror attempt to resynchronize the relationship. If the resync operation cannot find a common Snapshot copy, you must reestablish the relationship with the 'snapmirror initialize' command. For SnapVault, issue the command 'snapvault start -r' to make SnapVault attempt to resynchronize the relationship. If the 'snapvault start -r' command cannot find a common Snapshot copy, you must reestablish the relationship with the 'snapvault start' command. If you must reestablish a relationship, either run the baseline transfer to a new qtree, or remove the existing qtree before starting the transfer. Refer to the Data Protection Guide for details.

### Syslog Message

%s: destination transfer from %s:%s to %s:%s : base Snapshot copy for transfer no longer exists on the source.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination path.

**rplerror** (STRING): Replication error message.

**waferror** (STRING): WAFL® error.

## replication.dst.resync.failed

### Severity

ERROR

### Description

This message occurs when a SnapMirror® or SnapVault® resync operation is not successful.

### Corrective Action

See error message.

## Syslog Message

%s resync of %s to %s:%s : %s.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**dstpath** (STRING): Replication destination volume or qtree.

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source volume or qtree.

**error** (STRING): Error message.

## replication.dst.resync.success

### Severity

NOTICE

### Description

This message occurs when a SnapMirror® or SnapVault® resync operation is successful.

### Corrective Action

Operation successful; no corrective action required.

## Syslog Message

%s resync of %s to %s:%s was successful.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**dstpath** (STRING): Replication destination volume or qtree.

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source volume or qtree.

## replication.dst.setACLError

### Severity

ERROR

### Description

This event occurs when Data ONTAP encounters an illegal data stream during a logical backup operation. This is usually due to an OSSV primary software bug. Normally OSSV sets an ACL (Access Control List) on a regular file or directory. However, as a result of this bug, OSSV attempts to set an ACL on an Alternate Data Stream (ADS) of a file or directory. See bug 148879 for details.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**src\_base\_inum** (INT): The base inode number generated by the source. The source is trying to set an ACL on this inode, which happens to be a stream.

**dst\_base\_inum** (INT): The corresponding base inode number on the destination.

## replication.dst.setACLFailed

### Severity

ERROR

### Description

This message occurs when Data ONTAP® encounters an error while setting an ACL (Access Control List) for an inode on a SnapVault® or SnapMirror® destination. The replication transfer is aborted upon encountering this error.

### Corrective Action

(None).

### Syslog Message

Replication transfer from %s:%s to %s:%s: The destination failed to set ACL %d for source inode %d on destination inode %d. The replication transfer is aborted.

### Parameters

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination path.

**acl\_inode** (INT): Inode number of the source ACL.

**src\_base\_inum** (INT): Base inode number generated by the source. The destination is trying to set an ACL on this inode.

**dst\_base\_inum** (INT): Corresponding base inode number on the destination.

## replication.dst.throttleChange

### Severity

NOTICE

### Description

This event is issued when the replication source, receives a request from the destination to change the current throttle value.

### Corrective Action

(None).

### Syslog Message

Throttle value changed to %d Kbytes/s for destination %s:%s

### Parameters

**throttle** (INT): The new throttle value.

**dstfiler** (STRING): The SnapMirror destination filer

**dstpath** (STRING): The SnapMirror destination volume or qtree

## replication.dst.throttleChange.err

### Severity

NOTICE

## Description

This event is issued when the replication source, receives a request from the destination to change the current throttle value, but failed to complete the operation.

## Corrective Action

(None).

## Syslog Message

Failed to change throttle value to %d Kbytes/s for destination %s:%s : %s

## Parameters

**throttle** (INT): The new throttle value.  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination volume or qtree  
**error** (STRING): The error message

## replication.dst.wafErr

### Severity

ERROR

### Description

This message occurs when the system detects an error on a SnapMirror® or SnapVault® destination volume or qtree due to WAFL®.

### Corrective Action

(None).

### Syslog Message

%s: destination transfer from %s:%s to %s : %s:%s.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).  
**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source volume or qtree.  
**dstpath** (STRING): Replication destination volume or qtree.  
**rplerror** (STRING): Specific replication error message.  
**waferror** (STRING): Specific WAFL error message.

## replication.err events

### replication.err.snapCreateFailed

#### Severity

ERROR

#### Description

This message occurs when the replication engine fails to create an automatic Snapshot(tm) copy.

## Corrective Action

Make sure that the Snapshot limit is not reached for the volume. The current Data ONTAP® version allows a maximum of 255 Snapshot copies per volume.

## Syslog Message

%s transfer from %s:%s to %s:%s : %s.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination path.

**rplerror** (STRING): Replication error message.

**waferror** (STRING): WAFL® error message.

# replication.givebackcancel events

## replication.givebackCancel

### Severity

ERROR

### Description

This message occurs when a 'giveback' command is issued while a SnapMirror® or SnapVault® transfer is in progress.

### Corrective Action

(None).

### Syslog Message

%s transfer in progress or in-sync; canceling giveback.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

# replication.global events

## replication.global.throttle

### Severity

NOTICE

### Description

This event is issued when the request to change/set the throttling value for a replication association is failed. This is because the network layer does not support throttling traffic. This usually happens when the connection is going through a tcp offload device.

### Corrective Action

(None).

## Syslog Message

Failed to set global throttle value for transfer from %s:%s to %s:%s

## Parameters

**srcfiler** (STRING): The replication source filer  
**srcpath** (STRING): The replication src volume or qtree  
**dstfiler** (STRING): The replication destination filer  
**dstpath** (STRING): The replication destination volume or qtree

# replication.lcl events

## replication.lcl.noHostname

### Severity

ERROR

### Description

This message occurs when a SnapMirror® or SnapVault® transfer is started on a system without a host name.

### Corrective Action

Configure the hostname on the indicated system before attempting a SnapMirror or SnapVault transfer operation.

## Syslog Message

%s transfer could not be started because no host name is configured on this system.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on).  
**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source path.  
**dstfiler** (STRING): Replication destination system.  
**dstpath** (STRING): Replication destination path.  
**rplerror** (STRING): Replication error message.  
**waferror** (STRING): WAFL® error.

# replication.nomemory events

## replication.noMemory

### Severity

ALERT

### Description

This message occurs when there is insufficient memory for the replication transfer to continue.

### Corrective Action

The system is out of memory. This could be a transient situation. Reduce the load on the system or defer the transfer.

## Syslog Message

Replication transfer failed to allocate memory.

## Parameters

**file** (STRING): File in which the software error occurred.

**line** (INT): Line number in the file for the software error.

# replication.qtree events

## replication.qtree.dst.disk.error

### Severity

ALERT

### Description

This message occurs when the replication destination engine receives blocks with disk errors from the source for the first time in a transfer and continues with replication.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

SnapMirror: transfer from %s:%s to %s:%s: Source volume has disk errors. Continuing with replication. See /etc/log/snapmirror for inconsistency details.

## Parameters

**src\_file** (STRING): Source filer of this replication transfer.

**src\_path** (STRING): Source path of this replication transfer.

**dst\_file** (STRING): Destination filer of this replication transfer.

**dst\_path** (STRING): Destination path of this replication transfer.

## replication.qtree.dst.disk.error.abort

### Severity

ERROR

### Description

This message occurs when the replication destination engine receives blocks with disk errors from the source and aborts because of user-set options.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

SnapMirror: transfer from %s:%s to %s:%s: Source volume has disk errors. (source fid:%u,source fbn:%u,source snapid:%u). Aborting transfer.

## Parameters

**src\_file** (STRING): Source filer of this replication transfer.

**src\_path** (STRING): Source path of this replication transfer.

**dst\_file** (STRING): Destination filer of this replication transfer.

**dst\_path** (STRING): Destination path of this replication transfer.  
**sfileid** (INT): Identifying number for the suspect file  
**sfbn** (INT): File block number in which the disk error was found.  
**ssnapid** (INT): Incremental Snapshot(tm) copy of the source from which this file came.

## replication.qtree.dst.disk.error.summary

### Severity

ERROR

### Description

This message occurs when the replication involving blocks with disk errors happens on the replication destination.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

SnapMirror: transfer from %s:%s to %s:%s: Source volume has disk errors. %u blocks out of %u blocks replicated are inconsistent. See /etc/log/snapmirror for inconsistency details.

### Parameters

**src\_filer** (STRING): Source filer of this replication transfer.  
**src\_path** (STRING): Source path of this replication transfer.  
**dst\_filer** (STRING): Destination filer of this replication transfer.  
**dst\_path** (STRING): Destination path of this replication transfer.  
**total\_bad\_blks** (INT): Total number of blocks with disk errors received.  
**total\_blks** (INT): Total number of blocks received.

## replication.qtree.dst.inconsistent.acl

### Severity

ERROR

### Description

This message occurs when the replication destination engine receives an Access Control List (ACL) with disk errors from the source.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

SnapMirror: transfer from %s:%s to %s:%s: Source volume has disk errors. (source acl fid:%u,source snapid:%u). Aborting transfer.

### Parameters

**src\_filer** (STRING): Source filer of this replication transfer.  
**src\_path** (STRING): Source path of this replication transfer.  
**dst\_filer** (STRING): Destination filer of this replication transfer.  
**dst\_path** (STRING): Destination path of this replication transfer.  
**sfileid** (INT): Identifying number for the suspect ACL.  
**snapid** (INT): Incremental Snapshot(tm) copy of the source from which this ACL came.



## replication.qtree.dst.inconsistent.symlink

### Severity

ERROR

### Description

This message occurs when the replication destination engine receives a symbolic link with disk errors from the source.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

SnapMirror: transfer from %s:%s to %s:%s: Source volume has disk errors. (source symlink fid:%u,source snapid:%u). Aborting transfer.

### Parameters

**src\_filer** (STRING): Source filer of this replication transfer.

**src\_path** (STRING): Source path of this replication transfer.

**dst\_filer** (STRING): Destination filer of this replication transfer.

**dst\_path** (STRING): Destination path of this replication transfer.

**sfileid** (INT): Identifying number for the suspect symbolic link.

**snapid** (INT): Incremental Snapshot(tm) copy of the source from which this symbolic link came.

## replication.regoperr events

### replication.regOpErr

### Severity

ERROR

### Description

This message occurs when the system encounters an error during a registry operation related to SnapMirror® or SnapVault®.

### Corrective Action

The reasons for the error cited in the error message might indicate ways to address the problem.

### Syslog Message

Replication registry operation error: %s.

### Parameters

**reason** (STRING): Reason for registry operation failure.

## replication.resync events

### replication.resync.cmnsnpCnt

### Severity

NOTICE

## Description

This message occurs during a SnapMirror® or SnapVault® resync operation to inform the user of the number of common Snapshot(tm) copies detected between the source and destination.

## Corrective Action

(None).

## Syslog Message

Resynchronization of %s:%s to %s:%s found %d matching Snapshot copies; using the most recent one.

## Parameters

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination volume or qtree.

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source volume or qtree.

**cmnsnapcnt** (INT): Number of common Snapshot copies detected between the source and destination.

# replication.rmt events

## replication.rmt.noHostname

### Severity

ERROR

### Description

This message occurs when the remote system of a SnapMirror® or SnapVault® transfer does not have a host name.

### Corrective Action

Configure the host name on the remote system before attempting a SnapMirror or SnapVault transfer from that system.

### Syslog Message

%s transfer from %s:%s to %s:%s cannot be started because the remote system does not have a host name.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on).

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination path.

**rplerror** (STRING): Replication error message.

**waferror** (STRING): WAFL® error.

# replication.src events

## replication.src.err

**Severity**

ERROR

**Description**

This message occurs when the system detects an error on a SnapMirror® or SnapVault® source volume or qtree.

**Corrective Action**

(None).

**Syslog Message**

%s: source transfer from %s to %s:%s : %s.

**Parameters**

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**srcpath** (STRING): Replication source volume or qtree.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination volume or qtree.

**error** (STRING): Specific error message.

**replication.src.inconsistent.snapmap****Severity**

ERROR

**Description**

This message occurs when a replication source engine encounters a disk error in a snapmap block. Replication transfer will abort.

**Corrective Action**

Run wafiron as soon as possible, and then Contact NetApp technical support.

**Syslog Message**

WAFL metadata inconsistent due to disk errors in volume %s. The affected Snapshot copy ID is %d (level: %d). Aborting transfer.

**Parameters**

**volname** (STRING): Aggregate, traditional volume, or flexible volume name.

**snapid** (INT): Snapshot(tm) copy ID.

**level** (INT): Buffer level.

**replication.src.throttleChange****Severity**

NOTICE

**Description**

This event is issued when the replication source, receives a request from the destination to change the current throttle value.

**Corrective Action**

(None).

**Syslog Message**

Throttle value changed to %d Kbytes/s for destination %s:%s

**Parameters**

**throttle** (INT): The new throttle value.

**dstfiler** (STRING): The SnapMirror destination filer

**dstpath** (STRING): The SnapMirror destination volume or qtree

**replication.src.throttleChange.err****Severity**

ERROR

**Description**

This event is issued when the replication source, receives a request from the destination to change the current throttle value, but failed to complete the operation.

**Corrective Action**

(None).

**Syslog Message**

Failed to change throttle value to %d Kbytes/s for destination %s:%s : %s

**Parameters**

**throttle** (INT): The new throttle value.

**dstfiler** (STRING): The SnapMirror destination filer

**dstpath** (STRING): The SnapMirror destination volume or qtree

**error** (STRING): The error message

**replication.src.unknownCommand****Severity**

NOTICE

**Description**

This event is issued when the replication source, SnapMirror or SnapVault, receives a request from the destination which it does not understand. This can be seen when changes are made to the negotiation phase internally during development. This error message should NEVER show up on customer builds.

**Corrective Action**

This message is normally seen after upgrading the source or the destination. ONTAP version mismatch between the source and destination can cause this. One should not see this problem if the source and destination filers are running the same version of ONTAP.

**Syslog Message**

Replication source received an unknown command from the destination. This might be due to Data ONTAP version mismatch between the source and destination.

## Parameters

(None).

# replication.stats events

## replication.stats.notLoaded

### Severity

INFORMATIONAL

### Description

This message occurs during system boot when SnapMirror® or SnapVault® cannot load the statistics (for example, the number of transfers completed) from the registry. The statistics from the previous run of the system are lost and SnapMirror or SnapVault will start collecting the statistics from scratch.

### Corrective Action

(None).

### Syslog Message

%s could not load statistics from the registry.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

## replication.stats.notSaved

### Severity

INFORMATIONAL

### Description

This message occurs when SnapMirror® or SnapVault® cannot save some statistics (for example, the number of transfers completed) in the registry at the end of a transfer. The operation is retried at the end of the next SnapMirror or SnapVault transfer.

### Corrective Action

(None).

### Syslog Message

%s could not save statistics to the registry.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

# replication.status events

## replication.status.timeSkewed

### Severity

NOTICE

## Description

This message occurs when the time skew between the source and destination system is greater than 1 minute.

## Corrective Action

Use either the 'date' or 'rdate' command to synchronize the time between the systems.

## Syslog Message

%s: Clock on source %s is ahead of destination system %s.

## Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on.).

**srcfiler** (STRING): The host name of the source system.

**dstfiler** (STRING): The host name of the destination system.

# replication.upgrade events

## replication.upgrade.complete

### Severity

INFORMATIONAL

### Description

This message occurs when post-upgrade operations for SnapMirror® and SnapVault®, as a part of Data ONTAP upgrade, are complete.

### Corrective Action

(None).

### Syslog Message

Post-upgrade operations for SnapMirror and SnapVault are complete.

### Parameters

(None).

# replication.vol events

## replication.vol.activemapdelta.err

### Severity

ERROR

### Description

This message occurs when SnapMirror® or VolCopy replication encounters a disk error in the active map delta blocks on the destination and cannot continue replication.

### Corrective Action

For a SnapMirror update operation, break the SnapMirror relation and then resynchronize the relation. For a VolCopy or SnapMirror Initialize operation, destroy the destination volume and retry the operation. If the problem persists, contact NetApp technical support.

## Syslog Message

WAFL active map delta file is inconsistent due to disk errors '%s%s%s'. Aborting transfer.

## Parameters

**volname** (STRING): Flexible volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Volume unique identifier.

## replication.vol.dst.disk.error

### Severity

ERROR

### Description

This message occurs when the replication destination engine receives the first disk error from the replication source in a transfer. Replication will continue. Subsequent disk errors encountered during this transfer will not trigger this event, but will be logged to /etc/log/snapmirror. A replication.vol.dst.disk.error.summary event is logged at the end of the transfer to report on the total number of errors encountered.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

%s: %s: Source volume has disk errors. Continuing replication. See /etc/log/snapmirror for inconsistency details.

## Parameters

**rpltype** (STRING): Replication type - SnapMirror, Aggr copy, Vol copy.  
**transfer\_paths** (STRING): Replication source path and destination path.

## replication.vol.dst.disk.error.abort

### Severity

ERROR

### Description

This message occurs when the replication destination engine receives blocks with disk errors from the source. Replication will abort because the user selected the option to abort on disk error.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

%s: %s: Source volume has disk errors (vbn:%llu). Aborting transfer.

## Parameters

**rpltype** (STRING): Replication type - SnapMirror, Aggr copy, Vol copy  
**transfer\_paths** (STRING): Replication source path and destination path  
**vbn** (LONGINT): Volume Block Number

## replication.vol.dst.disk.error.summary

### Severity

ERROR

### Description

This message occurs on the replication destination in a transfer that encounters disk errors. Details about the errors are in /etc/log/snapmirror.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

%s: %s: Source volume has disk errors. %u blocks out of %llu blocks replicated are inconsistent. See /etc/log/snapmirror for inconsistency details.

### Parameters

**rpl\_type** (STRING): Replication type - SnapMirror, Aggr copy, Vol copy.  
**transfer\_path** (STRING): Replication source path and destination path.  
**total\_bad\_blocks** (INT): Total number of blocks with disk errors replicated.  
**total\_blocks** (LONGINT): Total number of blocks replicated.

## replication.vol.src.inconsistent.metadata

### Severity

ALERT

### Description

This message occurs when a replication source engine encounters a disk error in the file system metadata. Replication transfer will abort.

### Corrective Action

If the specified Snapshot copy IDs contain 0 (active file system), run wafiron to fix the active file system. To continue replication transfer, delete the specified Snapshot copies during the next reboot, either by using the boot menu command 'boot\_snap\_delete' or by using the CLI command 'vol online -r'. Then contact NetApp technical support.

### Syslog Message

%s: %s: WAFL metadata inconsistent due to disk errors. (vbn:%llu level:%u file\_type:%d) The affected snapshot IDs are %s. Aborting transfer.

### Parameters

**rpltype** (STRING): Replication type - SnapMirror, Aggr copy, Vol copy  
**transfer\_paths** (STRING): Replication source path and destination path.  
**vbn** (LONGINT): Virtual Volume Block Number  
**level** (INT): Buffer level.  
**file\_type** (INT): Inode type.  
**snapstring** (STRING): String containing the list of Snapshot copy IDs of the Snapshot copies that have the inconsistent blocks.



# replication.wormlog events

## replication.wormlog.appendFail

### Severity

ERROR

### Description

This event is issued when failed to append to the replication worm log file. This can be seen when the worm log volume is full, or other hardware failure in the log volume.

### Corrective Action

Check the free space and the sanity of the worm log volume. Add or replace hard disk for the worm log volume.

### Syslog Message

Failed to append to the replication worm log: %s.

### Parameters

**error\_msg** (STRING): Error message for the failure.

# replication.xfr events

## replication.xfr.checksum.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror®/SnapVault(tm) operation is not completed because a block received from the source system failed checksum verification. This might be caused by corruption in the network stream that went undetected by the TCP checksum.

### Corrective Action

Scheduled SnapMirror/SnapVault update operations are automatically retried and should succeed. However, if the problem persists, use the 'netstat -s' command to review the network statistics for TCP packets "discarded for bad checksums". If you find large numbers of such discards, diagnose the network between the source and destination systems for any problems. If TCP does not report any checksum errors and this problem still persists, contact NetApp technical support technical support.

### Syslog Message

%s transfer from %s:%s to %s:%s : %s.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on).

**srcfiler** (STRING): Replication source system.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination system.

**dstpath** (STRING): Replication destination path.

**rplerror** (STRING): Replication error message.

**waferror** (STRING): WAFL® error.

## replication.xfr.nonExistSnap

### Severity

NOTICE

### Description

A SnapMirror® or SnapVault(tm) operation was not completed because the specified snapshot could not be found. Any of the following scenarios can cause this event to occur: (1) The SnapVault schedules are configured incorrectly at the destination. (2) The 'snapmirror initialize' or 'snapmirror update' command was issued with the '-s' option, and the transfer cannot start. (3) The SnapVault or SnapMirror transfer cannot restart from a checkpoint.

### Corrective Action

For scheduled SnapVault updates, ensure that the snapshot name on the destination volume and source volume is the same. For 'snapmirror initialize' or 'update' commands from a named snapshot, ensure that the requested snapshot exists on the source volume. For restarted SnapVault or SnapMirror transfers, run the 'snapmirror abort' or 'snapvault abort' command with the '-h' option before retrying the transfer.

### Syslog Message

%s transfer from %s:%s to %s:%s : %s.

### Parameters

**rtype** (STRING): Replication type (SnapMirror, SnapVault, and so on).

**srcfiler** (STRING): Replication source appliance.

**srcpath** (STRING): Replication source path.

**dstfiler** (STRING): Replication destination appliance.

**dstpath** (STRING): Replication destination path.

**rplerror** (STRING): Replication error message.

**waferror** (STRING): WAFL® error.

# repository events

## repository.cleanup events

### Repository.Cleanup.BaseFile

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

ERROR

#### Description

This message occurs when there is a failure in the Infinite Volume while creating a stream on a nonexistent file, and the base file was left behind in the Vserver.

#### Corrective Action

Delete the file created from the mounted file system. Example: To delete a file from the Windows cmd shell: C:\Users\Administrator>del Z:\File.txt If the problem persists, contact NetApp technical support.

#### Syslog Message

Failure happened in the Infinite Volume while creating a stream on a nonexistent file, and the base file (handle: %s) was left behind in Vserver %u.

#### Parameters

**filehandle** (STRING): File handle of the base file that was created.

**vserverId** (INT): Vserver ID for which the error occurred.

## repository.create events

### Repository.Create.Placement.Failed

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

ERROR

#### Description

This message occurs during a create request from a client, when the Infinite Volume namespace redirector encounters one or more of the following: 1. It made several attempts to create the item on some data constituents, but failed. It is possible that some data constituents, or aggregates in the Infinite Volume, or nodes in the cluster are not in a healthy state. 2. It could not find a data constituent in which to create the requested item. It is possible that the Infinite Volume is misconfigured or the configuration state of the Infinite Volume is corrupted.

#### Corrective Action

Verify that the Infinite Volume is configured properly. Verify that all the constituents are online by using the 'volume show' command with the '-constituents' option. Verify that all the aggregates for the Infinite Volume are online by using the 'aggregate show' command. If the problem persists or you need assistance, contact

NetApp technical support.

### **Syslog Message**

A create request failed on the Infinite Volume. Reason: %s. This occurred in VServer %u.

### **Parameters**

**reason** (STRING): Reason for the failure.  
**vserverId** (INT): Vserver ID for which the error occurred.

## **repository.ctpupdate events**

### **Repository.CTPUpdate.Failed**

#### **Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### **Severity**

ERROR

#### **Description**

This message occurs when an update to the Create Time Placement (CTP) ruleset fails.

#### **Corrective Action**

Retry updating the ruleset by using the "vserver data-policy import" command. If the problem persists, contact NetApp technical support.

### **Syslog Message**

Create Time Placement (CTP) ruleset update failed on the Infinite Volume. Reason: %s.

### **Parameters**

**reason** (STRING): Reason for the failure.

## **repository.dangling events**

### **Repository.Dangling.Data.File**

#### **Deprecated**

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### **Severity**

ERROR

#### **Description**

This message occurs during a client file operation request, when the Infinite Volume namespace redirector runs into failures that prevent it from cleaning up a file on a data constituent. This dangling file is unnecessarily consuming resources.

#### **Corrective Action**

Contact NetApp technical support.

## Syslog Message

Failures caused the Infinite Volume namespace redirector to leave a dangling file (handle: %s) on a data constituent in Vserver %u.

## Parameters

**filehandle** (STRING): File handle of the file that is left dangling in the data constituent.  
**vserverId** (INT): Vserver ID for which the error occurred.

## Repository.Dangling.Redirector

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs during a client file operation request, when the Infinite Volume namespace redirector runs into a dangling redirector file on the namespace constituent.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Failures caused the Infinite Volume namespace redirector to leave a dangling redirector file %s on the namespace constituent in Vserver %u.

## Parameters

**filehandle** (STRING): File handle of the redirector file that is left dangling in the namespace constituent.  
**vserverId** (INT): Vserver ID for which the error occurred.

## repository.dc events

## Repository.DC.Dangling.Lock

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs during a client file operation request, when the Infinite Volume runs into failures that prevent it from cleaning up a transient or client-originated lock on the data constituent. If there are enough of these dangling locks on the data constituent, the lock resource might be exhausted, which causes certain operations to fail.

### Corrective Action

Clear the reported lock ID by using the 'vserver locks break' command at the advanced privilege level.

## Syslog Message

Failures prevented the cleanup of a lock with UUID %s on the data constituent in Vserver %u.

## Parameters

**lockid** (STRING): Lock UUID of the lock that is left dangling.

**vserverId** (INT): Vserver ID for which the error occurred.

# repository.inconsistent events

## Repository.Inconsistent.Attribute

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs during a client file operation request, when the Infinite Volume namespace redirector detects an inconsistency in the file attributes for a file in the namespace constituent.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

While servicing a protocol request, the Infinite Volume namespace redirector detected an inconsistency in the attributes of a file (handle: %s) in the namespace constituent in Vserver %u.

## Parameters

**filehandle** (STRING): File handle of the file for which the error occurred.

**vserverId** (INT): Vserver ID for which the error occurred.

# repository.invalid events

## Repository.Invalid.Redirector

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

NOTICE

### Description

This message occurs during a client file operation request, when the Infinite Volume namespace redirector detects an internal inconsistency in a redirector file in the namespace constituent.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

While servicing a protocol request, the Infinite Volume namespace redirector detected an internal inconsistency in a file (handle: %s) in the namespace constituent in Vserver %u.

## Parameters

**filehandle** (STRING): File handle of the file for which the error occurred.  
**vserverId** (INT): Vserver ID for which the error occurred.

# repository.ns events

## Repository.NS.Dangling.Lock

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs during a client file operation request, when the Infinite Volume namespace redirector runs into failures that prevent it from cleaning up a transient lock on the namespace constituent. If there are enough of these dangling locks on the namespace constituent, the lock resource might be exhausted, which causes certain operations to fail.

### Corrective Action

Clear the reported lock ID by using the 'vserver locks break' command at the advanced privilege level.

## Syslog Message

Failures prevented the cleanup of a transient lock with UUID %s on the namespace constituent in Vserver %u.

## Parameters

**lockid** (STRING): Lock UUID of the lock that is left dangling.  
**vserverId** (INT): Vserver ID for which the error occurred.

## Repository.NS.Dangling.Streams

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when a client wants to truncate the file and only the truncation of the data constituent file succeeds while the redirector truncation fails. This leaves behind dangling stream redirectors. These dangling stream redirectors cannot be read or modified and unnecessarily consume resources on the namespace constituent.

## Corrective Action

Go to controller system shell, and delete all streams corresponding to that redirector. Example of deleting dangling streams from the system shell: `bash-3.2$ ls /clus/vs0/NS/File.txt/nmatr%/ s1.txt s2.txt bash-3.2$ sudo rm -f /clus/vs0/NS/File.txt/nmatr%/*` If the problem persists, contact NetApp technical support.

## Syslog Message

Failures caused the Infinite Volume namespace redirector to leave dangling stream redirectors on file (handle: %s) on the namespace constituent in Vserver %u.

## Parameters

**filehandle** (STRING): File handle of the redirector base file that has dangling stream redirectors.  
**vserverId** (INT): Vserver ID for which the error occurred.

# repository.snapshot events

## Repository.Snapshot.Access

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when the Infinite Volume namespace redirector restores access to the Infinite Volume Snapshot(tm) copies after it recovers from an internal communication problem that had previously shut down access.

### Corrective Action

(None).

### Syslog Message

Infinite Volume namespace redirector restored access to the Snapshot copies after it recovered from an internal communication problem.

### Parameters

(None).

## Repository.Snapshot.NoAccess

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

ERROR

### Description

This message occurs when the Infinite Volume namespace redirector cannot get access to Snapshot(tm) mapping updates due to an internal communication issue, so it prevents access to the Infinite Volume Snapshot copies. Several scenarios can cause this condition, including the following: 1. The cluster administrator has made the node ineligible for the cluster. 2. Intense system loading can cause this problem



temporarily.

**Corrective Action**

If the cluster administrator has made the node ineligible for the cluster (through "cluster modify"), make the node eligible again. If the condition persists or occurs repeatedly, other serious issues might be the cause and you should contact NetApp technical support.

**Syslog Message**

Infinite Volume namespace redirector shut down access to the Snapshot copies due to an internal communication problem.

**Parameters**

(None).

# respmon events

## respmon.service events

### respMon.service.responsive

#### Severity

NOTICE

#### Description

This message occurs when a service being monitored is responding to requests.

#### Corrective Action

(None).

#### Syslog Message

A service (ID: %u) of type %s, running on node %s (ID: %u, Uuid: %s), is responsive.

#### Parameters

**serviceID** (INT): Unique identifier of the service (ServiceID).

**serviceType** (STRING): Service type.

**nodeName** (STRING): Name of the node on which the service is running.

**nodeID** (INT): Unique identifier of the node on which the service is running.

**nodeUuid** (STRING): Universal unique identifier of the node on which the service is running.

### respMon.service.unresponsive

#### Severity

NOTICE

#### Description

This message occurs when a service being monitored is unresponsive.

#### Corrective Action

(None).

#### Syslog Message

A service (ID: %u) of type %s, running on node %s (ID: %u, Uuid: %s), is unresponsive.

#### Parameters

**serviceID** (INT): Unique identifier of the service (ServiceID).

**serviceType** (STRING): Service type.

**nodeName** (STRING): Name of the node on which the service is running.

**nodeID** (INT): Unique identifier of the node on which the service is running.

**nodeUuid** (STRING): Universal unique identifier of the node on which the service is running.

# revert events

## revert.ufmt events

### revert.ufmt.cap.tier.full

#### Severity

ALERT

#### Description

This message occurs when the revert operation to convert the unified object format is paused because the FabricPool aggregate has reached the maximum supported size of the capacity tier. The revert scan will resume when additional space is available in the capacity tier.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Revert operation to the convert the object format from "unified" to "native" is paused on aggregate "%s" because the capacity tier usage has reached its maximum allowed size.

#### Parameters

**aggregate\_name** (STRING): Aggregate name.

### revert.ufmt.config.invalid

#### Severity

NOTICE

#### Description

This message occurs when the revert operation to convert the unified object format from "unified" to "native" is paused because the object store configuration associated with the FabricPool aggregate is not ready.

#### Corrective Action

This condition might be transient and the revert operation should resume automatically. If the problem persists, contact NetApp technical support.

#### Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on FabricPool aggregate "%s" because the associated object store configuration is not ready.

#### Parameters

**aggregate\_name** (STRING): Aggregate name.

### revert.ufmt.mirror.not.sync

#### Severity

NOTICE

## Description

This message occurs when the revert operation to convert the unified object format from "unified" to "native" is paused because the FabricPool mirror is not in sync with the primary object store attached to the FabricPool aggregate.

## Corrective Action

This condition might be transient, and the revert operation should resume automatically when the FabricPool mirror is synced with the attached primary object store. If the problem persists, contact NetApp technical support.

## Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on FabricPool aggregate "%s" because the FabricPool mirror is not in sync with the attached primary object store.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

## revert.ufmt.objstore.failed

### Severity

ERROR

### Description

This message occurs when the revert operation to convert the unified object format is paused because it is unable to read objects from an unresponsive object store. Object format revert operation will resume when the object store is responsive.

### Corrective Action

Verify that your intercluster LIF is online and functional by using the "network interface show" command. Also, check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. Additionally, ensure that the configuration of your object store has not changed, and that login and connectivity information is still accurate. If the condition persists, contact NetApp technical support.

### Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on aggregate "%s" because the object store is unavailable.

### Parameters

**aggregate\_name** (STRING): Aggregate name.

## revert.ufmt.perf.tier.full

### Severity

ALERT

### Description

This message occurs when the revert operation to convert the unified object format is paused because the FabricPool aggregate does not have enough free space to tier cold data. The revert scan will resume when additional space is available in the performance tier.

## Corrective Action

Either free space in the performance tier of the aggregate by moving volumes, removing volumes, or deleting data, or add space to the performance tier of the aggregate by adding disks. If the problem persists, Contact NetApp technical support.

## Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on aggregate "%s" because the current available space in the performance tier is below the minimum allowed size.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

## revert.ufmt.tlog.disabled

### Severity

NOTICE

### Description

This message occurs when the revert operation to convert the unified object format is paused because FabricPool tiering is temporarily disabled on a MetroCluster setup. The object format revert will resume when FabricPool tiering is enabled.

### Corrective Action

(None).

### Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on aggregate "%s" because FabricPool tiering is disabled.

### Parameters

**aggregate\_name** (STRING): Aggregate name.

## revert.ufmt.tlog.full

### Severity

NOTICE

### Description

This message occurs when the revert operation to convert the unified object format is paused because FabricPool tiering is temporarily suspended due to insufficient space to stage new objects. Tiering will resume when space is available after writing existing objects to the object store. The revert scan will resume when FabricPool tiering is enabled.

### Corrective Action

(None).

### Syslog Message

Revert operation to convert the object format from "unified" to "native" is paused on aggregate "%s" because FabricPool tiering is temporarily suspended on the aggregate.

**Parameters**

**aggregate\_name** (STRING): Aggregate name.

**revert.ufmt.tlog.punch.pause****Severity**

NOTICE

**Description**

This message occurs when the revert operation to convert unified object format is paused because the FabricPool aggregate has complete unified format objects waiting to be tiered to the object store. The revert operation will resume when FabricPool tiering has completed tiering all waiting "unified format" objects to the object store.

**Corrective Action**

(None).

**Syslog Message**

The operation to revert the object format to the native format paused on aggregate "%s" because the tiering of unified format objects to the object store is pending.

**Parameters**

**aggregate\_name** (STRING): Aggregate name.

# revertto events

## revertto.force events

### revertTo.force

**Severity**

NOTICE

**Description**

This message occurs when a revert-to process is being initiated with the "force" flag, and configuration errors have been detected.

**Corrective Action**

(None).

**Syslog Message**

Forcing revert in spite of detected configuration errors.

**Parameters**

(None).

## revertto.start events

### revertTo.start

**Severity**

NOTICE

**Description**

This message occurs when a revert to process is started.

**Corrective Action**

(None).

**Syslog Message**

Starting revert to %s.

**Parameters**

**target\_release** (STRING): Target release name to which the revert process was started.

## revertto.system events

### revertTo.system.reboot

**Severity**

NOTICE

**Description**

This message occurs when the system is rebooting to complete the revert process. During reboot, the RAID subsystem will be reverted. After the RAID subsystem reverts and the system reboots, the revert process is complete.

**Corrective Action**

(None).

**Syslog Message**

System reverting to work with Data ONTAP %s.

**Parameters**

**target\_release** (STRING): Target release name to which the revert process was started.



# rllib events

## rllib.duplicate events

### rllib.duplicate.ip.detected

#### Severity

ERROR

#### Description

This message occurs when an interconnect link IP address on this node has a conflict with an IP address on another node in the cluster.

#### Corrective Action

Reboot either this node or the other node reported in this event, while making sure that all interconnect links on the two nodes are connected, up, and operational.

#### Syslog Message

Interconnect link IP address %s on this node (System ID %d) is also in use by another node (System ID %d) in the cluster.

#### Parameters

**dupIP** (STRING): Conflicting IP address.  
**localSysid** (LONGINT): System ID of the local node.  
**remoteSysid** (LONGINT): System ID of the remote node.

## rllib.ifconfig events

### rllib.ifconfig.linkEvent

#### Severity

NOTICE

#### Description

This event is generated when there is a link event on the MCC/IP managed interface.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**ifname** (STRING): The name of the interface  
**eventType** (STRING): The event type

## rllib.mccip events

## **rllib.mccip.config.failure**

### **Severity**

ERROR

### **Description**

This event is generated if any error occurs while configuring MCC/IP.

### **Corrective Action**

(None).

### **Syslog Message**

Configuration error has occurred in "%s", while setting up "%s ",with value: "%s", (error string='%s', error code=%d).

### **Parameters**

**objectOper** (STRING): Operation being executed.

**objectType** (STRING): Type of Object - IP address/ Netmask/ Interface / Bootarg

**objectValue** (STRING): Value of IP address/ Netmask/ Interface / Bootarg

**errorString** (STRING): The error string

**errNum** (INT): The internal error code

## **rllib.mccip.config.missing**

### **Severity**

NOTICE

### **Description**

This event is generated if any of the environment data relating to the configuration of an MCC/IP DR port is missing on boot.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**objectOper** (STRING): Interface being configured.

## **rllib.remote events**

### **rllib.remote.node.event**

### **Severity**

NOTICE

### **Description**

This event is generated to print information about remote node such as IP address and sysid along with the event that occurred

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**peerIP** (STRING): IP address of remote node

**peerSysid** (LONGINT): Sysid of remote node

**localIP** (STRING): IP address of local node

**localPort** (STRING): Local rdma port name

**eventDesc** (STRING): Information about the event (Optional)

**rllib.remote.node.update****Severity**

NOTICE

**Description**

This event is generated when remote node sysid has changed as a part of headswap

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**oldSysid** (LONGINT): Old sysid of remote node

**sysid** (LONGINT): New sysid of remote node

# rlm events

## rlm.driver events

### rlm.driver.hourly.stats

#### Severity

ERROR

#### Description

This message occurs when the system encounters an error while trying to get hourly statistics from the Remote LAN Module (RLM).

#### Corrective Action

1. Issue the 'rlm status' command at the Data ONTAP® prompt to check whether the RLM is online. 2. If the RLM is online and this message persists, issue the 'rlm reboot' command at the Data ONTAP prompt to reboot the RLM.

#### Syslog Message

The software driver for the Remote LAN Module (RLM) detected a problem: %s.

#### Parameters

**reason** (STRING): This error could be due to a communication error with the RLM or because the RLM is not operational.

### rlm.driver.mailhost

#### Severity

ERROR

#### Description

This message occurs when Remote LAN module (RLM) setup verifies whether a mailhost specified in ONTAP can be reached. In this case, RLM setup cannot connect to the specified mailhost.

#### Corrective Action

1. Verify that a valid mailhost is configured in Data ONTAP by checking the system AutoSupport configuration. 2. Ensure that Data ONTAP can successfully connect to the specified mailhost by invoking test command to invoke AutoSupport.

#### Syslog Message

RLM setup could not access the mailhost specified in Data ONTAP.

#### Parameters

(None).

### rlm.driver.network.failure

#### Severity

ERROR

## Description

This message occurs when the system encounters a failure during network configuration of the Remote LAN Module (RLM). The system cannot assign the RLM a DHCP (Dynamic Host Configuration Protocol) or fixed IP address.

## Corrective Action

1. Check whether the network cable is correctly plugged into the RLM network port.
2. Check the link status LED on the RLM.
3. The RLM supports a 10/100 Ethernet network in autonegotiation mode. Verify that the network that the RLM is connected to has autonegotiation enabled.

## Syslog Message

The network configuration of the Remote LAN Module (RLM) failed due to cable or network errors.

## Parameters

(None).

## rlm.driver.timeout

### Severity

ERROR

### Description

This message occurs when there is a failure during communication with the Remote LAN Module (RLM) firmware.

### Corrective Action

1. Issue the 'rlm status' command at the Data ONTAP® prompt to check whether the RLM is online.
2. If the RLM is operational and this message persists, issue the 'rlm reboot' command at the Data ONTAP prompt to reboot the RLM. After the reboot, this message should no longer occur. If the message occurs again, contact NetApp technical support and explain that you already performed the preceding steps.

### Syslog Message

The software driver for the Remote LAN Module (RLM) detected a problem: %s.

### Parameters

**reason** (STRING): This error could be due to a communication error with the RLM or because the RLM is not operational.

## rlm.firmware events

### rlm.firmware.upgrade.reqd

### Severity

NOTICE

### Description

This message occurs when the Remote LAN Module (RLM) firmware version and the Data ONTAP® software version are incompatible and cannot communicate correctly about a particular capability.

## Corrective Action

For more information about upgrading your firmware, check your support provider's web site or knowledgebase. After the RLM is running the new firmware, this message should no longer occur. If the problem persists or you need further assistance, contact NetApp technical support.

## Syslog Message

The RLM firmware %s is incompatible with Data ONTAP for %s.

## Parameters

**current\_rlm\_fw\_version** (STRING): Current firmware version on the RLM.

**capability** (STRING): RLM capability needed by Data ONTAP.

## rlm.firmware.version.unsupported

### Severity

NOTICE

### Description

This message occurs when the firmware on the Remote LAN Module (RLM) is not an official release and must be upgraded.

## Corrective Action

For more information about upgrading your firmware, check your support provider's web site or knowledgebase. After the RLM is running the new firmware, this message should no longer occur. If the problem persists or you need further assistance, contact NetApp technical support.

## Syslog Message

The RLM firmware version %s is not an official release.

## Parameters

**current\_rlm\_fw\_version** (STRING): Current firmware version on the RLM.

## rlm.heartbeat events

### rlm.heartbeat.resumed

### Severity

INFORMATIONAL

### Description

This event is generated when the appliance detects resumption of Remote LAN Module (RLM) heartbeat notifications indicating that the RLM is now available. The earlier issue indicated by the rlm.heartbeat.stopped event has been resolved.

## Corrective Action

(None).

## Syslog Message

Heartbeat received from the Remote LAN Module (RLM).

## Parameters

(None).

## rlm.heartbeat.stopped

### Severity

ERROR

### Description

This event is generated when the filer does not receive Remote LAN Module (RLM) heartbeat notifications. The RLM card and the filer exchange heartbeat messages so that they can detect when one or the other is unavailable. This event is generated when the filer has not received an expected heartbeat message from the RLM.

### Corrective Action

Connect to the RLM CLI and issue 1. rlm version 2. rlm config 3. priv set advanced 3.1 rlm log debug 3.2 rlm log messages 4. Run RLM system diagnostics 5. Consult RLM Troubleshooting Guide

### Syslog Message

Have not received a heartbeat from the Remote LAN Module (RLM) in last %d seconds.

### Parameters

**max\_anx** (INT): Heartbeats have not been received from the Remote Lan Module (RLM) for this period of time in seconds.

## rlm.network events

### rlm.network.link.down

#### Severity

ERROR

#### Description

This message occurs when the Remote LAN Module (RLM) detects a link error on the RLM network port. This can happen if a network cable is not plugged into the RLM network port. It can also happen if the network that the RLM is connected to cannot run at 10/100 Mbps.

#### Corrective Action

1. Check whether the network cable is correctly plugged into the RLM network port.
2. Check the link status LED on the RLM.
3. The RLM supports a 10/100 Mbps Ethernet network in autonegotiation mode. Verify that the network that the RLM is connected to supports autonegotiation to 10/100 Mbps or is running at one of those speeds; otherwise, RLM network connectivity won't work.

#### Syslog Message

Remote LAN Module (RLM) network port link down due to cable or network errors.

#### Parameters

(None).

# rlm.notconfigured events

## rlm.notConfigured

### Severity

ERROR

### Description

This message occurs weekly to remind you to configure the Remote LAN Module (RLM). The RLM is a physical device that is incorporated into your system to provide remote access and remote management capabilities. To use the full functionality of RLM, you need to configure it first.

### Corrective Action

1. Use the 'rlm setup' command to configure the RLM. If necessary, use the 'rlm status' command to obtain its MAC address. 2. Use the 'rlm status' command to verify the RLM network configuration. 3. Use the 'rlm test autosupport' command to verify that the RLM can send AutoSupport email. Note that AutoSupport mailhosts and recipients must be properly configured in Data ONTAP before issuing this command.

### Syslog Message

The system's Remote LAN Module (RLM) is not configured. Use the 'rlm setup' command to configure it.

### Parameters

(None).

# rlm.orftp events

## rlm.orftp.failed

### Severity

NOTICE

### Description

A communication error occurred while sending or receiving information from the Remote Lan Module (RLM).

### Corrective Action

1. Issue the 'rlm status' command at the Data ONTAP prompt to check whether the RLM is operational. 2. If the RLM is operational and this message persists, issue the 'rlm reboot' command at the Data ONTAP(tm) prompt to reboot the RLM. 3. If this message persists after the RLM has been rebooted, contact NetApp technical support.

### Syslog Message

RLM communication error, %s.

### Parameters

**reason** (STRING): This error could be due to : 1. Communication error while the information is being sent or received. 2. Remote Lan Module (RLM) being non-operational.



# rlm.snmp events

## rlm.snmp.traps.off

### Severity

INFORMATIONAL

### Description

This message occurs each time a system boots, if the advanced privilege level in Data ONTAP® was used to disable the SNMP Trap feature of the Remote LAN Module (RLM). This message also occurs when the SNMP Trap capability was disabled and a user invokes a Data ONTAP command to use the RLM to send an SNMP trap.

### Corrective Action

RLM SNMP Trap support is currently disabled. To enable this feature, set the 'rlm.snmp.traps' option to On.

### Syslog Message

RLM's SNMP Trap support is currently disabled. Set option 'rlm.snmp.traps' to 'On' to enable this feature.

### Parameters

(None).

# rlm.ssh events

## rlm.ssh.access.resetup

### Severity

ERROR

### Description

RLM SSH access list needs to be resetup.

### Corrective Action

Use 'system sp\_ssh\_access\_list modify' CLI to resetup RLM SSH access list.

### Syslog Message

RLM SSH access list needs to be resetup.

### Parameters

(None).

# rlm.userlist events

## rlm.userlist.update.failed

### Severity

ERROR

## Description

There was an error while updating user information for the Remote Lan Module (RLM). When user information is updated on Data ONTAP®, the RLM is also updated with the new changes. This allows users to login to the RLM.

## Corrective Action

1. Issue the 'rlm status' command at the Data ONTAP() prompt to check whether the Remote Lan Module (RLM) is operational.
2. If the RLM is operational, and this message persists, issue the 'rlm reboot' command at the Data ONTAP® prompt to reboot the RLM.
3. Retry the operation that caused the error message.
4. If this message persists after the RLM has been rebooted, contact NetApp technical support.

## Syslog Message

Error updating RLM user information, %s.

## Parameters

**reason** (STRING): User information update for the Remote Lan Module (RLM) may have failed due to : 1. Communication error with the RLM. 2. RLM may not be operational.

# rpc events

## rpc.unspecified events

### rpc.unspecified.error

#### Deprecated

Deprecated as of version 9.2 where this RPC mechanism was withdrawn.

#### Severity

NOTICE

#### Description

This event is issued when turning an ypgroup update fails to occur.

#### Corrective Action

(None).

#### Syslog Message

RPC: Failed (unspecified error)

#### Parameters

(None).

## rpc.xprt events

### rpc.xprt.limit

#### Deprecated

Deprecated as of version 9.2 where this RPC mechanism was withdrawn.

#### Severity

ERROR

#### Description

This message occurs when the number of active connections to a side band Remote Procedure Call (RPC) reaches the limit. This limit is not configurable.

#### Corrective Action

To check for stale client connections, use the output from the CLI "netstat" command.

#### Syslog Message

Active TCP connections to side band RPC service with program number %u has reached a limit of %d connections.

#### Parameters

**prog** (INT): TCP program number of RPC service to which the connection request arrived.

**max\_conn** (INT): Maximum number of active connections to RPC service.

## rpc.xprt.vfiler.limit

### Deprecated

Deprecated as of version 9.2 where this RPC mechanism was withdrawn.

### Severity

ERROR

### Description

This message occurs when the number of active connections to side band Remote Procedure Call (RPC) services of a vFiler unit reaches the limit. This limit is not configurable.

### Corrective Action

To check for stale client connections, use the output from the CLI "netstat" command.

### Syslog Message

Active TCP connections to side band RPC services of vFiler %s has reached a limit of %d connections.

### Parameters

**vfiler** (STRING): Name of the vFiler unit to which the connection request arrived.

**vf\_max\_conn** (INT): Maximum number of active connections to sideband RPC services of this vFiler unit.

# rpl events

## rpl.check events

### rpl.check.unusable.log.file

**Severity**

ERROR

**Description**

This message occurs when SnapMirror® Fixer detects inconsistency in the verifier log file generated by a previous SnapMirror 'check' operation. The SnapMirror 'fixer' operation will be terminated.

**Corrective Action**

Run the SnapMirror 'fixer' command.

**Syslog Message**

The SnapMirror verifier log is inconsistent and could not be used.

**Parameters**

(None).

## rpl.fixer events

### rpl.fixer.launch.autocheck

**Severity**

NOTICE

**Description**

This message occurs when SnapMirror® Fixer cannot use the log files generated by a previous SnapMirror 'check' operation. A new SnapMirror 'check' operation will be started.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror Fixer could not use the %s log file. A new SnapMirror 'check' operation will be started.

**Parameters**

**filetype** (STRING): Type of SnapMirror checker log file.

## rpl.snapshot events

### rpl.snapshot.delete.error

**Severity**

ERROR

## Description

This message occurs when SnapMirror® or SnapVault® cannot delete one of its Snapshot(tm) copies.

## Corrective Action

SnapMirror or SnapVault tries to delete this Snapshot copy during the next update. However, if no further updates are possible or planned, you should delete this Snapshot copy manually.

## Syslog Message

(None).

## Parameters

**snapname** (STRING): Snapshot copy name.

**volname** (STRING): Volume name.

**error** (STRING): WAFL® error message.

# rpl.src events

## rpl.src.lun.old\_clone

### Severity

ERROR

### Description

This message occurs when the replication source finds a LUN clone in a Snapshot(tm) copy that was created on a Data ONTAP® version earlier than 7.3 and was attempting to transfer as a LUN clone instead of a LUN. Transferring a LUN clone in a Snapshot(tm) copy that was created on Data ONTAP version earlier than 7.3 as a LUN clone might lead to divergence on the replication destination.

### Corrective Action

Only SnapVault® transfers that are managed using Ontapi can get into this situation. If SnapVault® restore is using a named Snapshot copy transfer, then choose a different named Snapshot copy. If SnapVault® is not using named Snapshot copy transfer, then delete the Snapshot copy and try again.

### Syslog Message

%s transfer from %s:%s to %s:%s is attempting to transfer a LUN clone from a pre Data ONTAP 7.3 Snapshot copy %s to the destination as a LUN clone.

### Parameters

**transfer\_type** (STRING): Type of the transfer.

**source\_filer** (STRING): System name of the replication source.

**source\_path** (STRING): Path of the qtree where the clones reside.

**dest\_filer** (STRING): System name of the replication destination.

**dest\_path** (STRING): Replication destination path.

**snapname** (STRING): Name of the backup Snapshot copy of the LUN clone.

# rpt events

## rpt.disk events

### rpt.disk.uuid.gen.failed

**Severity**

ERROR

**Description**

This message occurs when the generation of a disk UUID fails.

**Corrective Action**

Investigate the Directed Attached (DA) disk to fix the error, and then reboot the device.

**Syslog Message**

Disk name %s and disk unit number %d of RAID passthru with error message %s.

**Parameters**

**d\_name** (STRING): Disk name.

**d\_unit** (INT): Unit number.

**errmsg** (STRING): Whether the device failed or is misconfigured.

# rsh events

## rsh.enabled events

### rsh.enabled

**Severity**

NOTICE

**Description**

This message occurs when rsh is enabled in the cluster using the (privilege: advanced) "security protocol" command.

**Corrective Action**

(None).

**Syslog Message**

Insecure protocol service rsh has been enabled in the cluster.

**Parameters**

(None).

## rsh.firewall events

### rsh.firewall.policy.remove

**Severity**

NOTICE

**Description**

This message occurs when rsh is automatically removed from firewall policies during a revert operation from Data ONTAP® 9.0 to 8.3. Rsh is removed because it is an insecure protocol and it is by default disabled in Data ONTAP 9.0. In Data ONTAP 8.3, the presence of rsh in firewall policies after a revert operation enables rsh. To prevent this, the rsh firewall policy is removed during the operation.

**Corrective Action**

Use the "system services firewall policy create" command to reenable rsh.

**Syslog Message**

Rsh removed from firewall policy because of Data ONTAP revert operations. Rsh is now disabled.

**Parameters**

(None).

## rsh.rcmd events

### rsh.rcmd.port.bind



**Severity**

ERROR

**Description**

This message occurs when there is an error in binding a reserved port.

**Corrective Action**

(None).

**Syslog Message**

rcmd.c: Bind error %d on reserved port %d.

**Parameters**

**errNum** (INT): Error number.

**port** (INT): Port number.

## rsh.rcmd.port.used

**Severity**

ERROR

**Description**

This message occurs when there is no port available for binding to a socket, because all the reserved ports are in use.

**Corrective Action**

(None).

**Syslog Message**

rcmd.c: All reserved ports have been used.

**Parameters**

(None).

## rsh.setsockopt events

### rsh.setsockopt.failure

**Severity**

ERROR

**Description**

This message occurs when the setsockopt option on the RemoteShell(RSH) socket fails; RSH is not available.

**Corrective Action**

Use the error message to determine the problem.

**Syslog Message**

RSH setsockopt error: option level '%s', option value '%s', error code '%d', error message '%s'.

## Parameters

**socketlevel** (STRING): Socket option level parameter.  
**socketoption** (STRING): Socket option value parameter.  
**errorcode** (INT): Error code returned from setsockopt.  
**reason** (STRING): Error message for the failure.

# rsh.socket events

## rsh.socket.failure

### Severity

ERROR

### Description

The rshd daemon could not create a socket.

### Corrective Action

Disable and enable rshd daemon using options rsh.enable.

### Syslog Message

RSH socket error: %s.

## Parameters

**reason** (STRING): The reason why the rshd daemon could not create a socket.

# rshd events

## rshd.failure events

### rshd.failure

#### Severity

ERROR

#### Description

This event is generated after rshd failed to start.

#### Corrective Action

Restart the rsh daemon for the affected vfiler by using the vfiler run command to set the rsh.enable option to off and then to on for that vfiler.

#### Syslog Message

rshd failed to start for vfiler %s

#### Parameters

**vfilername** (STRING): The name of the vfiler that attempted to start rshd

# s3 events

## s3.bucket events

### s3.bucket.invalid.accessKey

#### Severity

ERROR

#### Description

This message occurs when the access key is invalid and it has been used to access the S3 bucket/object.

#### Corrective Action

Verify that the access key is valid and it has been configured correctly before accessing the bucket/object.

#### Syslog Message

S3 %s operation failed while accessing bucket "%s" on S3 server %s because access key "%s" is invalid.

#### Parameters

**operation** (STRING): S3 request name.

**bucket** (STRING): Name of the S3 bucket.

**serverHostname** (STRING): Hostname of the S3 server.

**accessKey** (STRING): Access key.

### s3.bucket.notFound

#### Deprecated

This event is renamed to "osc.bucket.notFound"

#### Severity

ERROR

#### Description

This message occurs when the bucket is not found while trying to access an S3 object.

#### Corrective Action

Verify that the bucket has been created on the S3 server and the bucket name has been correctly configured in ONTAP before accessing the bucket.

#### Syslog Message

Bucket "%s" not found on S3 server %s when performing %s operation.

#### Parameters

**bucket** (STRING): Name of the S3 bucket.

**serverHostname** (STRING): Hostname of the S3 server.

**operation** (STRING): S3 request name.

### s3.bucket.signatureMismatch

**Deprecated**

This event is renamed to "osc.signatureMismatch"

**Severity**

ERROR

**Description**

This message occurs when the request signature sent to the S3 server does not match the signature calculated by the server. This could be due to either an incorrect secret key or an internal ONTAP software error in calculating the signature.

**Corrective Action**

Verify that the secret access key is configured correctly. If it is correct, contact NetApp technical support for assistance.

**Syslog Message**

S3 %s operation server-calculated request signature does not match the signature sent to S3 server %s for bucket "%s".

**Parameters**

**operation** (STRING): S3 request name.

**serverHostname** (STRING): Hostname of the S3 server.

**bucket** (STRING): Name of the S3 bucket.

# saml events

## saml.cfg events

### saml.cfg.idp.download

#### Severity

NOTICE

#### Description

This message occurs when the Security Assertion Markup Language (SAML) Service Provider (SP) configuration fails to download the Identity Provider (IdP) metadata. Any partial configuration is removed automatically on this node only.

#### Corrective Action

Verify that the IdP metadata URI is reachable via the clustered Data ONTAP® "network ping" or "network ping6" command. If you need further assistance, contact NetApp technical support.

#### Syslog Message

SAML SP IdP metadata download from %s was not completed successfully. Error: '%s'.

#### Parameters

**idp\_uri** (STRING): IdP metadata location.

**reason** (STRING): Reason for the download failure.

### saml.cfg.setup.incomplete

#### Severity

NOTICE

#### Description

This message occurs when the Security Assertion Markup Language (SAML) Service Provider (SP) configuration is not completed successfully. Any partial configuration is removed automatically on this node only. This failure does not prevent authentication of any user. However, the Identity Provider (IdP) cannot retrieve the SP metadata from this node. This node automatically attempts to complete the configuration when the 'security saml modify -is-enabled true' command is executed.

#### Corrective Action

Use the 'security saml status show' command to see the status of all nodes. Verify the health of all nodes. To create the configuration without enabling SAML authentication, use the 'security saml delete' and 'security saml create' commands to re-create the configuration. If this behavior persists, contact NetApp technical support.

#### Syslog Message

SAML SP configuration creation was not completed successfully.

#### Parameters

(None).

## **saml.cfg.sync.fail**

### **Severity**

ERROR

### **Description**

This message occurs when the Security Assertion Markup Language (SAML) Service Provider (SP) configuration fails to synchronize on this node. This synchronization is most commonly performed due to SAML SP configuration changes or as new nodes come online. This failure might impact SAML-based authentication of any user, but will not impact data access.

### **Corrective Action**

Use the 'security saml status show' command to see the status of all nodes. Verify the health of all nodes. If this behavior persists, contact NetApp technical support.

### **Syslog Message**

SAML SP configuration modification was not completed successfully.

### **Parameters**

(None).

# san events

## san.changedhomenode events

### san.ChangedHomeNode

**Severity**

NOTICE

**Description**

This message occurs when the home node of an existing SAN LIF is modified during LIF placement.

**Corrective Action**

(None).

**Syslog Message**

Home node changed from %s to %s while placing LIF %s in Vserver %s.

**Parameters**

**old\_node\_name** (STRING): Name of the old home node of the existing SAN LIF.

**node\_name** (STRING): Name of the new home node of the existing SAN LIF.

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

## san.changedhomeport events

### san.ChangedHomePort

**Severity**

NOTICE

**Description**

This message occurs when the home port of an existing SAN LIF is modified during LIF placement.

**Corrective Action**

(None).

**Syslog Message**

Home port changed from %s to %s while placing LIF %s in Vserver %s.

**Parameters**

**old\_port\_name** (STRING): Name of the old home port of the existing SAN LIF.

**port\_name** (STRING): Name of the new home port of the existing SAN LIF.

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

## san.duplicateipaddress events



## san.DuplicateIpAddress

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations of clustered ONTAP®, when an IP LIF is created in the source cluster, but one of the storage virtual machines (SVMs) in the destination cluster is already configured with an IP LIF having the same IP address.

### Corrective Action

Provide a unique IP address for the primary SVM IP LIF when creating or modifying the LIF in the source cluster. After modifying the IP address, check that SVM configuration state is healthy by using the "metrocluster vserver show -vserver <vserver>" command. If Vserver configuration state is not healthy, resynchronize the SVM with its partner by using the "metrocluster vserver resync" command.

### Syslog Message

One of the SVM's in the destination cluster is configured with an %s LIF with the same IP address %s. Provide a different unique IP address for the primary SVM %s LIF %s when creating or modifying the LIF in the source cluster.

### Parameters

**destination\_data\_protocol** (STRING): Data protocol of the destination storage virtual machine (SVM) LIF.  
**IpAddress** (STRING): IP address existing in the destination cluster.  
**primary\_data\_protocol** (STRING): Data protocol of the primary SVM LIF.  
**lif** (STRING): Name of the primary SVM IP LIF.

## san.duplicatewwpn events

### san.DuplicateWWPN

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations of clustered Data ONTAP®, when an FC LIF is created in the source cluster, but one of the Vservers in the destination cluster is already configured with an FC LIF having the same WWPN.

### Corrective Action

Provide a unique WWPN for the primary Vserver FC LIF when creating or modifying the LIF in the source cluster.

### Syslog Message

One of the Vservers in the destination cluster is configured with an FC LIF with the same WWPN %s. Provide a different unique WWPN for the primary Vserver FC LIF %s when creating or modifying the LIF in the source cluster.

### Parameters

**wwpn** (STRING): WWPN existing in the destination cluster.  
**lif** (STRING): Name of the primary Vserver FC LIF.

# san.fcadapter events

## san.fcadapter.init.failure

### Severity

ERROR

### Description

This message occurs when FC adapters hosting LIFs fails to initialize during node boot.

### Corrective Action

Reboot the affected node. If the problem persists, contact NetApp technical support.

### Syslog Message

The following FC adapters with configured LIFs failed to initialize: %s.

### Parameters

**adapters** (STRING): List of FC adapters that failed to initialize.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

# san.fcnvmeunsupportedport events

## san.FcNvmeUnsupportedPort

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP has an FC NVMe (Nonvolatile Memory Express) logical interface (LIF) hosted on an unsupported FC port.

### Corrective Action

Use the "fcp adapter show -fields data-protocols-supported" command to view the list of valid Fibre-Channel ports in the fabric. Identify the Fibre-Channel ports that supports FC-NVMe LIFs with the help of the "data-protocols-supported" field. Bring down the FC-NVMe LIF using the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -status-admin down" command and move the FC-NVMe LIF to an FC port which support FC-NVMe LIFs using the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -home-port <FC\_NVME\_LIF\_supported\_FC\_port\_name> -home-node <Node\_name\_of\_supported\_FC\_port>" command. Bring the FC-NVMe LIF back up using the "network interface modify -vserver <vserver\_name> -lif <lif\_name> -status-admin up" command, run the "metrocluster check run" command and then verify that the LIF placement succeeded by using the "metrocluster check lif show -vserver <vserver\_name> -lif <lif\_name>" command.

### Syslog Message

Port "%s" on node "%s" does not support hosting FC-NVMe LIF "%s".

## Parameters

**port** (STRING): Unsupported port that hosts the FC-NVMe LIF.  
**node** (STRING): Name of the node on which the LIF is hosted.  
**lif** (STRING): Name of the FC-NVMe LIF hosted on unsupported FC port.

# san.initiator events

## san.initiator.login

### Severity

NOTICE

### Description

This message occurs when ONTAP detects that an initiator which was previously not logged on is now logged on.

### Corrective Action

(None).

### Syslog Message

Missing login for initiator "%s", Vserver "%s", portset "%s", has been logged on.

## Parameters

**initiator** (STRING): Initiator that was previously not logged on.  
**vserver** (STRING): Name of the Vserver.  
**portset** (STRING): Name of the portset that the initiator must be logged in to.  
**vserver\_uuid** (STRING): UUID of the Vserver.

## san.initiator.login.missing

### Severity

NOTICE

### Description

This message occurs when ONTAP detects that an initiator configured in one or more igroups is not logged in, which is required to ensure full SAN connectivity. Access to LUNs might be disrupted if a failover occurs.

### Corrective Action

Verify the host and fabric configuration to ensure connectivity to the node.

### Syslog Message

Initiator "%s" is not logged in to any LIFs on this node for Vserver "%s", portset "%s". Access to LUNs might be disrupted if a failover occurs.

## Parameters

**initiator** (STRING): Initiator that is not logged in.  
**vserver** (STRING): Name of the Vserver.  
**portset** (STRING): Name of the portset that the initiator must be logged in to.  
**vserver\_uuid** (STRING): UUID of the Vserver.

## san.initiator.login.not.req

### Severity

NOTICE

### Description

This message occurs when ONTAP detects that an initiator which was previously not logged on no longer needs to be logged on.

### Corrective Action

(None).

### Syslog Message

Missing login for initiator "%s", Vserver "%s", portset "%s" is no longer required.

### Parameters

**initiator** (STRING): The initiator that was previously not logged on.

**vserver** (STRING): Name of the Vserver.

**portset** (STRING): Name of the portset that the initiator was not logged on to.

**vserver\_uuid** (STRING): UUID of the Vserver.

## san.iscsilifrepaired events

### san.iscsiLifRepaired

### Severity

NOTICE

### Description

This message occurs when an iSCSI LIF is repaired during Vserver initialization. The LIF did not have a relative target port ID (rtpid) assigned, so a unique ID was associated with this LIF.

### Corrective Action

None.

### Syslog Message

The iSCSI LIF "%s" on Vserver "%s" was repaired during Vserver initialization. The LIF did not have a relative target port ID (rtpid) assigned, so a unique ID was associated with this LIF.

### Parameters

**lif\_name** (STRING): lgroup name.

**vserver\_name** (STRING): Vserver name.

## san.lifmove events

### san.lifmove.fail

### Severity

EMERGENCY

## Description

This message occurs when a LIF move failed and the recovery action to restore the LIF has also failed. LUNs are no longer accessible through this LIF. The LIF must be recovered manually.

## Corrective Action

Restore the LIF using the 'network interface create' command. Also, restore any iscsi access list settings and FQDN settings using 'iscsi interface' commands. Specific recovery commands are listed in the 'recover\_commands' parameter. For more information or assistance, contact NetApp technical support.

## Syslog Message

Moving LIF %s of vservers %s from %s:%s to %s:%s failed. Manual recovery required.

## Parameters

**lif\_name** (STRING): LIF name.

**vservers** (STRING): Vserver name.

**src\_node** (STRING): LIF was moving from this node.

**src\_port** (STRING): LIF was moving from this port.

**dest\_node** (STRING): LIF was moving to this node.

**dest\_port** (STRING): LIF was moving to this port.

**address** (STRING): LIF address.

**recovery\_commands** (STRING): Commands to recover the LIF and associated settings.

# san.lifmoveafterswitchback events

## san.lifmoveafterswitchback

### Severity

ERROR

### Description

This message occurs during MetroCluster(tm) switchback, when a LIF cannot be placed on one of the disaster recovery (DR) nodes because it is offline. The LIF will be placed on the HA partner of the DR node.

### Corrective Action

After the MetroCluster switchback operation is completed successfully, bring online the MetroCluster DR node. Change the home node of the LIF to the MetroCluster DR node by using the "network interface modify" command.

### Syslog Message

MetroCluster DR node %s is offline. The LIF %s belonging to Vserver %s will be placed on the HA partner node %s.

### Parameters

**dr\_node\_name** (STRING): Name of the MetroCluster DR node that is offline.

**lif** (STRING): LIF name.

**vservers** (STRING): Vserver name.

**node\_name** (STRING): Name of the HA partner node of the MetroCluster DR node.

# san.lifrepfieldsmissing events

## san.LifRepFieldsMissing

### Severity

ERROR

### Description

This message occurs when the system cannot replicate a SAN logical interface (LIF) in a MetroCluster(tm) configuration, because one or more attributes of the source LIF cannot be determined.

### Corrective Action

Verify that the home node of the LIF created or modified on the primary cluster is up. Also verify that the LIF's operational status is up, by using the "network interface show -vserver <vserver\_name> -lif <lif\_name>-fields status-oper" command. Check whether the Vserver configuration state is healthy by using the "metrocluster vserver show -vserver <vserver\_name>" command. If the configuration state is unhealthy, resynchronize the Vserver with its partner by using the "metrocluster vserver resync" command.

### Syslog Message

Attempted LIF replication failed on the destination cluster node %s for SAN LIF %s for Vserver %s.

### Parameters

**node** (STRING): Name of the destination cluster node (DR partner) that attempted LIF replication.

**lif** (STRING): Name of the SAN LIF.

**vserver** (STRING): Name of the Vserver.

# san.lifspernodelimit events

## san.LifsPerNodeLimit

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP cannot replicate a SAN or NVMe logical interface (LIF) in the destination cluster due to the LIF limits per node per Vserver.

### Corrective Action

Make sure all the nodes are online by running "storage failover show" command on both the clusters. If a node is taken-over by its HA partner, giveback the node using "storage failover giveback -ofnode <name\_of\_node\_to\_giveback>". Once the nodes are online and healthy, run the "metrocluster vserver resync -vserver <vserver\_name>" from the source cluster. Verify the vserver's configuration-state is healthy using "metrocluster vserver show -vserver <vserver\_name>" command. Contact technical support for assistance if configuration-state of the vserver does not change to healthy.

### Syslog Message

Metrocluster does not support more than %s %s LIFs per node per Vserver.

## Parameters

**limit** (STRING): LIF limit supported on Metrocluster

**data\_protocol** (STRING): Data protocol of the LIF that exceeds the LIF per node per Vserver limit.

# san.mccpurgeverifyfailed events

## san.MCCPurgeVerifyFailed

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered Data ONTAP® systems when a SAN cache purge verify operation fails. The corresponding Vservers are not available to take part in the subsequent switchover or switchback operation until the purge verify operation succeeds.

### Corrective Action

Contact NetApp technical support for assistance in purging the SAN cache.

### Syslog Message

The following Vservers: %s failed in the SAN cache purge verify operation. Contact NetApp technical support for assistance in purging the SAN cache.

### Parameters

**vserver** (STRING): List of Vservers that failed the purge verify operation.

# san.mccvserverinitfailed events

## san.MCCVserverInitFailed

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered Data ONTAP® systems when SAN Vserver failed to initialize during a switchover or switchback operation.

### Corrective Action

Contact NetApp technical support for assistance in initialization of the SAN Vserver.

### Syslog Message

The Vserver %s failed to initialize. Reason: %s. Contact NetApp technical support for assistance in initialization of the Vserver.

### Parameters

**vserver** (STRING): The SAN Vserver that failed to initialize.

**reason** (STRING): The reason for failure.

## san.mccvserverinitpending events

### san.MCCVserverInitPending

#### Severity

ALERT

#### Description

This message occurs in MetroCluster™ configurations on clustered Data ONTAP® systems when SAN Vservers exceed the time allotted for initialization during a switchover or switchback operation. This is not a failure, but a delay in initialization.

#### Corrective Action

Contact NetApp technical support for assistance in analyzing the reason for delay in Vserver initialization.

#### Syslog Message

The following Vservers: %s exceeded the time allotted (%u seconds) for Vserver initialization. Contact NetApp technical support to analyze the reason for delay in Vserver initialization.

#### Parameters

**pending\_vservers** (STRING): List of the SAN Vservers pending Vserver initialization.

**san\_budget\_in\_msecs** (INT): Maximum time (in seconds) the system waits for the SAN Vservers to complete initialization during a switchover or switchback operation.

## san.missinglicense events

### san.MissingLicense

#### Severity

ERROR

#### Description

This message occurs during replication of the FC/iSCSI service configuration, when the required service is not licensed or has expired in the cluster.

#### Corrective Action

Add the missing or expired license in the cluster by using the "license add" command.

#### Syslog Message

%s license could be missing or expired in the cluster.

#### Parameters

**license\_name** (STRING): Name of the missing license.

## san.noconnectedfcport events

### san.NoConnectedFCport



## Severity

ERROR

## Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP cannot replicate an FC/FC-NVMe (Nonvolatile Memory Express) logical interface (LIF) to a candidate port in the destination cluster. A candidate port is one that satisfies the following conditions: 1) It must reside on the disaster recovery (DR) partner node. 2) It must be connected to the same FC fabric as the home port of the primary Vserver FC/FC-NVMe LIF.

## Corrective Action

Check the cabling to ensure that ports on the destination cluster node (DR partner) are connected to the same fabric as the source cluster port hosting the primary Vserver FC/FC-NVMe LIF. Use the "fcp adapter show -fields fabric-name" command to view the port-fabric connectivity information. After ensuring proper fabric connectivity, run the "metrocluster check lif repair-placement" command on the source cluster with the Vserver name and LIF name as inputs to perform lif placement followed by the "metrocluster check lif show" command to verify the results of the command.

## Syslog Message

There are no ports on node "%s" connected to fabric "%s" in the destination cluster. Cannot place "%s" LIF "%s" hosted on the source cluster port %s of %s adapter type for Vserver "%s".

## Parameters

**node** (STRING): Name of the destination cluster node (DR partner) that does not have a port connected to the fabric.

**fabric\_name** (STRING): Name of the fabric that the source cluster port hosting the primary Vserver FC/FC-NVMe LIF is connected to.

**data\_protocol** (STRING): Data Protocol of the FC/FC-NVMe LIF.

**lif** (STRING): Name of the FC/FC-NVMe LIF to be placed.

**port** (STRING): Source cluster port that hosts the primary Vserver FC/FC-NVMe LIF.

**adapter\_type** (STRING): Adapter type of the source cluster port that hosts the primary Vserver FC/FC-NVMe LIF.

**vserver** (STRING): Name of the secondary Vserver in the MetroCluster relationship.

## san.NoConnectedFCport.prim

## Severity

ERROR

## Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when the source cluster FC port that is hosting an FC/FC-NVMe LIF of the primary Vserver is not connected to any fabric. The LIF cannot be placed correctly at the destination cluster because fabric information is not available.

## Corrective Action

Check the wiring of the source cluster FC port to ensure that it is connected to the fabric. Use the "fcp adapter show -fields fabric-name" command to view the port-fabric connectivity information. After ensuring proper fabric connectivity, enter the "metrocluster check lif repair-placement -vserver <vserver\_name> -lif <lif\_name>" command in the source cluster to initiate the LIF placement at the secondary cluster. After LIF placement has been completed, enter the "metrocluster check lif show -vserver <vserver\_name> -lif <lif\_name>" to see whether lif placement succeeded or not.

## Syslog Message

The source cluster FC port %s that is hosting "%s" LIF %s on primary Vserver %s is not connected to any fabric.

## Parameters

**home\_port** (STRING): Name of the source cluster FC port.

**data\_protocol** (STRING): Data Protocol of the lif.

**lif** (STRING): Name of the FC/FC-NVMe LIF.

**vserver** (STRING): Name of the primary Vserver.

# san.noconnectediscsiport events

## san.NoConnectediSCSIPort

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered ONTAP® systems, when ONTAP cannot replicate an IP logical interface (LIF) to a candidate port in the destination cluster. A candidate port is one that both resides on the disaster recovery (DR) partner node and is connected to the same subnet as the home port of the primary SVM IP LIF.

### Corrective Action

Check the wiring of the network ports and ensure that there are ports on the destination cluster node (DR partner) that are connected to the same subnet as the source cluster port hosting the primary SVM IP LIF. If the wiring is correct but this failure is still reported when you run the "metrocluster check lif show" command, ensure that the source LIF is operationally up. If the admin-status of the source LIF is down, bring it up using the "network interface modify" command. When the source LIF is up, run the "metrocluster check lif repair-placement" command from the source cluster, with the SVM name and the LIF name as inputs. Enter the "metrocluster check lif show" command again to check whether LIF placement succeeded or not.

## Syslog Message

Subnet discovery failed during LIF replication. There might not be any ports on the destination cluster node %s that are connected to subnet %s in IP space %s to place the %s LIF %s for SVM %s.

## Parameters

**node** (STRING): Name of the destination cluster node (DR partner) that does not have a port connected to the subnet.

**subnet** (STRING): The subnet that the primary cluster port hosting the primary SVM IP LIF is connected to.

**ipspace\_name** (STRING): Name of the IP space that the primary cluster port hosting the primary SVM IP LIF is connected to.

**data\_protocol** (STRING): Data protocol of the LIF.

**lif** (STRING): Name of the IP LIF.

**vserver** (STRING): Name of the secondary SVM in the MetroCluster relationship.

# san.nofcportsonnode events

## san.NoFCPortsOnNode

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered Data ONTAP® systems when an attempt to place an FCP LIF on a particular node fails because the system cannot find any active ports on that node.

### Corrective Action

List the offline FC ports on the DR partner node by using the "fcp adapter show -node <node\_name;> -state down" command. Modify the state of the offline FC ports with the following command: "fcp adapter modify -node <node\_name;> -adapter <adapter\_name;> -state up". Once the FC ports in question are online, resynchronize the Vserver configuration using the "metrocluster vserver resync" command. With the command "fcp adapter show -node <node\_name;>" verify that FC ports are present. If the FC ports are absent in the destination cluster node (DR partner), contact customer support to provision the FC ports.

### Syslog Message

Cannot find online FC ports on Node "%s" for placing LIF "%s" in Vserver "%s".

### Parameters

**node\_name** (STRING): Node name.

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

## san.nofcportwithadapertype events

### san.NoFCportwithAdapterType

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP cannot replicate an FC/FC-NVMe logical interface (LIF) to a candidate port in the destination cluster. A candidate port is one that satisfies the following conditions: 1) It must reside on the disaster recovery (DR) partner node. 2) It must be of the same adapter type (CNA/FC) as the home port of the primary Vserver FC/FC-NVMe LIF.

### Corrective Action

Ensure that there are ports in the destination cluster node (DR partner) that have the same adapter type (FC/CNA) as the source cluster port that hosts the primary Vserver FC/FC-NVMe LIF. Use the "fcp adapter show -fields physical-protocol" command to view the FC port-adapter type information. Run the "metrocluster check lif repair-placement" command in the source cluster with the Vserver name and LIF name as inputs after ensuring that there are ports in the destination cluster node (DR partner) that have the same adapter type. After running the command mentioned previously, run the "metrocluster check lif show" to see whether lif placement succeeded or not.

### Syslog Message

There are no ports on node "%s" in the destination cluster that have the adapter type "%s" required to host the "%s" LIF "%s" for Vserver "%s".

## Parameters

- node** (STRING): Name of the destination cluster node (DR partner) that does not have a port of the same adapter type as the port in the source cluster.
- adapter\_type** (STRING): Adapter type of the source cluster port that hosts the primary Vserver FC/FC-NVMe LIF.
- data\_protocol** (STRING): Data Protocol of the lif.
- lif** (STRING): Name of the FC/FC-NVMe LIF to be placed.
- vserver** (STRING): Name of the secondary Vserver in the MetroCluster relationship.

# san.noifgrpport events

## san.NoIfGrpPort

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered ONTAP® systems, when ONTAP places a LIF on a physical port because an ifgrp port cannot be found on the destination cluster node (DR partner) to host an IP LIF.

### Corrective Action

Ensure that ifgrp port exists; otherwise create an ifgrp port on the DR partner by taking the following steps. Run the "network port show" command to view the available ports on the DR partner. Use the "network port ifgrp create" command to create ifgrp port and "network port ifgrp add-port" command to populate an ifgrp port. After ifgrp port is available on the DR partner, run the "metrocluster check lif repair-placement" command in the source cluster with the SVM name and LIF name as inputs . After running this command, use the "metrocluster check lif show" command to see whether LIF placement succeeded or not.

### Syslog Message

There are no ifgrp ports on node "%s" in IPspace "%s" to host the %s LIF "%s" for SVM "%s".

## Parameters

- node** (STRING): Name of the DR partner that does not have an ifgrp port.
- ipspace\_name** (STRING): Name of the IPspace.
- data\_protocol** (STRING): Data protocol of the LIF.
- lif** (STRING): Name of the IP LIF.
- vserver** (STRING): Name of the SVM for which the IP LIF was incorrectly placed.

# san.nonetworkportadapertype events

## san.NoNetworkPortAdapterType

### Severity

ERROR

### Description

This message occurs in MetroCluster™ configurations on clustered ONTAP® systems, when ONTAP cannot find a port on the disaster recovery (DR) node that is of the same adapter type (converged network adapter (CNA)/network interface card (NIC)) as the home port of the primary storage virtual machine (SVM) data LIF.

## Corrective Action

Ensure that there are ports in the destination cluster node (DR partner) that have the same adapter type as the source cluster port that hosts the primary SVM IP LIF. Use the "net port show" and "run -node <node-name> sysconfig -ca" commands to view the network port adapter type information. After ensuring same adapter type port is available on the DR partner, run the "metrocluster check lif repair-placement" command in the source cluster with the SVM name and LIF name as inputs after ensuring that there are ports in the destination cluster node (DR partner) that have the same adapter type. After running the repair placement command, run the "metrocluster check lif show" command to see whether LIF placement succeeded or not.

## Syslog Message

There are no ports on node "%s" in the destination cluster that have the adapter type "%s" in IPspace "%s" required to host the secondary SVM "%s" %s LIF "%s".

## Parameters

**node** (STRING): Name of the destination cluster node (DR partner) that does not have a port of the same adapter type as the port in the source cluster.

**adapter\_type** (STRING): Adapter type of the source cluster port.

**ipspace\_name** (STRING): Name of the IPspace.

**vserver\_name** (STRING): storage virtual machine (SVM) that owns the LIF.

**data\_protocol** (STRING): Data protocol of the LIF.

**lif** (STRING): Name of the IP LIF.

# san.noportsinipospace events

## san.NoPortsInIpspace

### Severity

ERROR

### Description

This message occurs when an attempt to place an IP LIF on a particular node fails because the system cannot find any ports in an IPspace on that node.

### Corrective Action

Add at least one port belonging to the node to the IPspace by using the "network port broadcast-domain add-ports" command and then resynchronize the SVM configuration using the "metrocluster vserver resync" command.

### Syslog Message

Cannot find ports in IPspace "%s" on node "%s" for placing LIF "%s" in SVM "%s".

### Parameters

**ipspace\_name** (STRING): IPspace name.

**node\_name** (STRING): Node name.

**lif** (STRING): LIF name.

**vserver\_name** (STRING): SVM name.

# san.nosamenamedfcport events

## san.NoSameNamedFCport

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP cannot find the same named port to place an FC/FC-NVMe logical interface (LIF) in the destination cluster. A same named port is an FC port in the destination cluster that has the same name as that of the FC port in the source cluster hosting the FC/FC-NVMe logical interface (LIF).

### Corrective Action

Check the admin state of the destination cluster FC port to ensure that it is up. Use the "fcp adapter show -fields status-admin" command to view the admin state of the target FC ports. After ensuring the same named port is up, enter the "metrocluster check lif repair-placement -vserver <vserver\_name> -lif <lif\_name>" command in the source cluster to initiate the LIF placement at the secondary cluster. After LIF placement has been completed, enter the "metrocluster check lif show -vserver <vserver\_name> -lif <lif\_name>" to see whether LIF placement succeeded or not.

### Syslog Message

Same named FC port %s as the one in the source cluster that is hosting %s LIF %s on primary Vserver %s is not found in the destination cluster.

### Parameters

**home\_port** (STRING): Name of the source cluster FC port.

**data\_protocol** (STRING): Data Protocol of the lif.

**lif** (STRING): Name of the FC/FC-NVMe LIF.

**vserver** (STRING): Name of the primary Vserver.

## san.nosamenamediscsiport events

### san.NoSameNamedIscsiPort

### Severity

ERROR

### Description

This message occurs when a MetroCluster™ configuration cannot be implemented, because the destination cluster for an IP logical interface (LIF) does not appear to contain a port with the same name as the one used within the source cluster.

### Corrective Action

Check whether the broadcast domain of the destination cluster network port belongs to the IPspace configured for the destination SVM by using the "network ipspace show -vservers <vserver\_name> -ports <same\_named\_network\_port\_name>" command. Verify that the destination cluster network port is up by using the "network port show -fields link" command. After the destination cluster network port is up, initiate the LIF placement at the destination cluster by using the "metrocluster check lif repair-placement -vserver <vserver\_name> -lif <lif\_name>" command in the source cluster. Verify that the LIF placement succeeded by using the "metrocluster check lif show -vserver <vserver\_name> -lif <lif\_name>" command.

## Syslog Message

Destination cluster does not contain a port with the same name as network port %s in the source cluster that is hosting %s LIF %s on primary SVM %s.

## Parameters

**home\_port** (STRING): Name of the source cluster network port.

**data\_protocol** (STRING): Data protocol of the LIF.

**lif** (STRING): Name of the IP LIF.

**vserver** (STRING): Name of the primary SVM.

# san.nosupportedssamenamedfcport events

## san.NoSupportedSameNamedFCport

### Severity

ERROR

### Description

This message occurs in MetroCluster(tm) configurations on clustered Data ONTAP® systems, when Data ONTAP cannot find the same named port that support the FCP/FC-NVMe data-protocol to place an FC logical interface (LIF) in the destination cluster. A same named port is an FC port in the destination cluster that has the same name as that of the FC port in the source cluster hosting the FC logical interface (LIF).

### Corrective Action

Before enabling the Non-Shared Fabrics lif placement mechanism, please ensure that the same-named ports in the corresponding DR-partner node at the destination cluster have the same capabilities as the source cluster node's port.

## Syslog Message

Same named FC port %s as the one in the source cluster that is hosting LIF %s on primary Vserver %s, does not support the %s data-protocol in the destination cluster.

## Parameters

**home\_port** (STRING): Name of the destination cluster FC port.

**lif** (STRING): Name of the FC/FC-NVMe LIF.

**vserver** (STRING): Name of the primary Vserver.

**data\_protocol** (STRING): Data-Protocol of the LIF.

# san.novlanport events

## san.NoVlanPort

### Severity

ERROR

### Description

This message occurs when a clustered Data ONTAP® system is in a MetroCluster™ configuration. The system cannot find a virtual LAN (VLAN) port on the destination cluster node (that is, the DR partner) to host an IP LIF, and has instead chosen to place the LIF on a physical port.

## Corrective Action

Ensure that the destination cluster node (DR partner) has ports connected to the same VLANs as in the source with the command "network port show -node <node> -fields vlan-tag". Ensure that the DR partner has VLAN interfaces similar to those in the source with the command "network port show -node <node> -fields type". If necessary, create a VLAN interface or interfaces on the DR partner node using the "network port vlan create" command. After ensuring that the proper VLAN port exists, run the "metrocluster check lif repair-placement" command on the source cluster, with the SVM name and LIF name as inputs to perform LIF placement. Then verify the results with the command "metrocluster check lif show".

## Syslog Message

There are no VLAN ports with tag "%s" on node "%s" in IPspace "%s" to host the %s LIF "%s" for SVM "%s".

## Parameters

**node** (STRING): Name of the DR partner that does not have a VLAN port.

**vlan\_tag** (STRING): VLAN tag.

**ipspace\_name** (STRING): Name of the IPspace.

**data\_protocol** (STRING): Data protocol of the LIF.

**lif** (STRING): Name of the IP LIF.

**vserver** (STRING): Name of the SVM for which the IP LIF was incorrectly placed.

# san.optimizedpersonality events

## san.optimizedPersonality.mismatch

### Severity

ERROR

### Description

This message occurs when a mismatch is detected between the san-optimized personality of a node and the cluster.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

The san-optimized value "%s" does not match the san-optimized value of the cluster.

### Parameters

**san\_optimized** (STRING): The san-optimized value for the node.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

# san.slmSupportCheckFailed events

## san.slmSupportCheckFailed

### Severity

ERROR



## Description

This message occurs during reboot when SAN capability is enabled, the cluster has more than 8 nodes and lun.reporting.nodes options is disabled. It leads to host instability as the number of paths exposed exceeds the threshold that the host stacks can handle.

## Corrective Action

Set the option "lun.reporting\_nodes.enable" to "on" with the command "vserver options lun.reporting\_nodes.enable on" and enable Selective LUN Map on existing maps with the command "lun mapping add-reporting-nodes -vserver \* -path \* -igroup \* -local-nodes".

## Syslog Message

Clusters with more than 8 nodes must have SLM support enabled.

## Parameters

(None).

# san.volrehostigroupcreatefailed events

## san.volRehostIgroupCreateFailed

### Severity

ERROR

### Description

This message occurs on clustered Data ONTAP® systems, during Volume Rehost operation when Data ONTAP fails to create igroup on the destination Vserver.

### Corrective Action

Please create the igroup using "igroup create".

### Syslog Message

Failed to create igroup "%s" for Vserver "%s". Reason: "%s".

### Parameters

**igroup\_name** (STRING): Igroup name.  
**vserver\_name** (STRING): Vserver name.  
**reason** (STRING): Reason for failure.

# san.volrehostlunmapcreated events

## san.volRehostLunMapCreated

### Severity

INFORMATIONAL

### Description

This message occurs on clustered Data ONTAP® systems for Volume Rehost operation on successful creation of the LUN map the destination Vserver.

### Corrective Action

(None).

### Syslog Message

LUN "%s" successfully mapped with igroup "%s" for Vserver "%s".

### Parameters

**lun\_path** (STRING): LUN path.

**igroup\_name** (STRING): Igroup name.

**vserver\_name** (STRING): Vserver name.

## san.volrehostlunmapfailed events

### san.volRehostLunMapFailed

#### Severity

ERROR

#### Description

This message occurs on clustered Data ONTAP® systems, during Volume Rehost operation when Data ONTAP fails to create LUN map the destination Vserver.

#### Corrective Action

Use "lun map create" command to create the LUN map.

#### Syslog Message

Failed to map LUN "%s" with igroup "%s" for Vserver "%s". Reason: "%s".

#### Parameters

**lun\_path** (STRING): LUN path.

**igroup\_name** (STRING): Igroup name.

**vserver\_name** (STRING): Vserver name.

**reason** (STRING): Reason for failure.

## san.volrehostsanbestpractice events

### san.volRehostSANBestPractice

#### Severity

ERROR

#### Description

This message occurs on clustered Data ONTAP® systems for Volume Rehost operation when the destination Vserver does not comply with the SAN best practice.

#### Corrective Action

Please check if there are SAN LIFS on the volume owning node and its HA partner

**Syslog Message**

Node "%s" for Vserver "%s" does not comply SAN Best Practice to create LUN maps.

**Parameters**

**node\_name** (STRING): Node name.

**vserver\_name** (STRING): Vserver name.

# sas events

## sas.8001 events

### sas.8001.firmware.fault

**Severity**

NOTICE

**Description**

This message occurs when the PM8001 serial-attached SCSI (SAS) host bus adapter reports an adapter firmware error.

**Corrective Action**

(None).

**Syslog Message**

Detected firmware error on serial-attached SCSI (SAS) adapter %s. Scratchpad registers: 0x%x, 0x%x, 0x%x, 0x%x.

**Parameters**

**adapterName** (STRING): Name of the associated serial-attached SCSI (SAS) host bus adapter.

**scratchpad0** (INHEX): Contents of the host bus adapter scratchpad register 0.

**scratchpad1** (INHEX): Contents of the host bus adapter scratchpad register 1.

**scratchpad2** (INHEX): Contents of the host bus adapter scratchpad register 2.

**scratchpad3** (INHEX): Contents of the host bus adapter scratchpad register 3.

## sas.8072 events

### sas.8072.firmware.fault

**Severity**

NOTICE

**Description**

This message occurs when the PM8072 serial-att (SAS) host bus adapter reports an adapter firmware error.

**Corrective Action**

(None).

**Syslog Message**

Detected firmware error on serial-attached SCSI (SAS) adapter %s. Scratchpad register: 0x%x.

**Parameters**

**adapterName** (STRING): Name of the associated serial-attached SCSI (SAS) host bus adapter.

**sctratchpad1** (INHEX): Contents of the host bus adapter scratchpad register.

# **sas.adapter.events**

## **sas.adapter.bad.WWN**

### **Severity**

ERROR

### **Description**

The base WWN associated with the SAS adapter is of invalid format. A correct WWN will be synthesized. The adapter will continue to function but this error indicates a manufacturing problem.

### **Corrective Action**

The controller should be replaced with a correctly programmed one.

### **Syslog Message**

The SAS WWN "%s" associated with adapter %s is bad. A valid one will be synthesized.

### **Parameters**

**WWN** (STRING): The bad WWN.

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## **sas.adapter.debug**

### **Severity**

INFORMATIONAL

### **Description**

The SAS adapter driver debug event.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**debug\_string** (STRING): The debug statement that occurred.

## **sas.adapter.failed**

### **Severity**

ERROR

### **Description**

The SAS adapter driver could not recover the adapter after resetting it multiple times. The adapter will be put offline.

### **Corrective Action**

If the adapter is in use check cabling and if connected to disk shelves check the seating of IOM cards and disks. If the problem keeps occurring try replacing the adapter. Finally if the issue is still not resolved please

contact NetApp technical support for assistance in diagnosing a SAS issue.

### **Syslog Message**

Failed to recover SAS %s after multiple resets. Offlining the adapter.

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## **sas.adapter.firmware.download**

### **Severity**

INFORMATIONAL

### **Description**

The firmware is being updated on the SAS adapter.

### **Corrective Action**

(None).

### **Syslog Message**

Updating firmware on SAS adapter %s from version %s to version %s.

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**oldFirmwareVersion** (STRING): The version of the firmware currently on the adapter.

**newFirmwareVersion** (STRING): The version of the firmware being downloaded to the adapter.

## **sas.adapter.firmware.fault**

### **Severity**

NOTICE

### **Description**

A firmware fault was detected on the SAS adapter and it is being reset to recover.

### **Corrective Action**

(None).

### **Syslog Message**

Detected firmware fault 0x%x on SAS adapter %s.

### **Parameters**

**faultCode** (INHEX): The firmware fault code.

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## **sas.adapter.firmware.update.failed**

### **Severity**

EMERGENCY

**Description**

The firmware on the SAS adapter was unable to be updated.

**Corrective Action**

The adapter should be replaced as soon as possible. The SAS adapter driver will attempt to continue using the adapter without updating the firmware image.

**Syslog Message**

Unable to update firmware on SAS adapter %s.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**sas.adapter.not.ready****Severity**

ERROR

**Description**

The SAS adapter did not become ready after being reset.

**Corrective Action**

The SAS adapter driver will automatically attempt to recover from this error. If the error keeps occurring the adapter may need to be replaced.

**Syslog Message**

SAS adapter %s did not become ready.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**sas.adapter.offline****Severity**

INFORMATIONAL

**Description**

The SAS adapter is now offline.

**Corrective Action**

(None).

**Syslog Message**

SAS adapter %s is now offline.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**sas.adapter.offlining**

**Severity**

INFORMATIONAL

**Description**

The SAS adapter is going offline after all outstanding I/O requests have finished.

**Corrective Action**

(None).

**Syslog Message**

Offlining SAS adapter %s.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**sas.adapter.online****Severity**

INFORMATIONAL

**Description**

The SAS adapter is now online.

**Corrective Action**

(None).

**Syslog Message**

SAS adapter %s is now online.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**sas.adapter.online.failed****Severity**

ERROR

**Description**

The SAS adapter failed to come online.

**Corrective Action**

If the adapter is in use check cabling and if connected to disk shelves check the seating of IOM cards.

**Syslog Message**

SAS adapter %s failed to come online.

**Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.



## **sas.adapter.onlining**

### **Severity**

INFORMATIONAL

### **Description**

The SAS adapter is in the process of going online.

### **Corrective Action**

(None).

### **Syslog Message**

Onlining SAS adapter %s.

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## **sas.adapter.unexpected.status**

### **Severity**

NOTICE

### **Description**

The SAS adapter returned an unexpected status and will be reset to recover.

### **Corrective Action**

(None).

### **Syslog Message**

Unexpected status 0x%x received from SAS adapter %s.

### **Parameters**

**status** (INHEX): The status returned from the adapter.

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## **sas.adapter.wwn.update.fail**

### **Severity**

ERROR

### **Description**

This message occurs when the WWN, a unique world wide name that identifies the adapter, fails to update. This message only applies to internal lab systems.

### **Corrective Action**

Contact your administrator.

### **Syslog Message**

SAS adapter "%s" fails to update WWN "%s".

## Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter.

**WWN** (STRING): WWN of the associated SAS host bus adapter.

## **sas.adapter.wwn.updated**

### Severity

NOTICE

### Description

This message occurs when an SAS host bus adapter's worldwide name (WWN) is updated successfully. The WWN is a unique name that identifies the serial-attached SCSI (SAS) host bus adapter.

### Corrective Action

(None).

### Syslog Message

SAS adapter "%s" WWN updated from "%s" to "%s"

## Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter.

**OldWWN** (STRING): Old WWN of the associated SAS host bus adapter.

**WWN** (STRING): New WWN of the associated SAS host bus adapter.

## **sas.cable events**

### **sas.cable.degraded**

#### Severity

ERROR

#### Description

This message occurs when the cable attached to SAS (Serial Attached SCSI) adapter port is functioning in a degraded mode.

#### Corrective Action

This condition implies one or more of the lanes in a 4-wide cable is functioning at a less-than-optimal speed, or is not functioning and has been disabled. The port will continue functioning but in a degraded, lower-performing mode. The cable should be replaced at some point in the future.

#### Syslog Message

Cable attached to SAS port "%s" is functioning in a degraded mode.

#### Parameters

**adapterName** (STRING): The name of the associated SAS host bus adapter.

### **sas.cable.linkratelow**

#### Severity

NOTICE

## Description

This message occurs when the expander or cable attached to the SAS (serial-attached SCSI) host bus adapter port is functioning at a speed lower than its specification. This could happen because the HBA is connected to an IOM that supports a lower speed, or because of a speed degradation due to a faulty component, or because of a lower common shelf stack speed.

## Corrective Action

(None).

## Syslog Message

Expander or cable attached to SAS port %s is operating at %s, which is lower than expected.

## Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter.

**currentSpeed** (STRING): Current link rate for SAS host bus adapter, cable, and IOM expander.

## **sas.cable.unqualified**

### Severity

EMERGENCY

### Description

This message occurs when a Serial Attached SCSI (SAS) cable attached to a controller port has not been qualified by NetApp. Unqualified cables might exhibit unpredictable behavior or experience higher error rates than qualified cables.

### Corrective Action

Replace the cable with one qualified by NetApp

### Syslog Message

Cable attached to SAS port %s has not been qualified by NetApp

### Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter port.

## **sas.config events**

### **sas.config.bad**

#### Severity

ERROR

#### Description

This message occurs when the SAS (Serial Attached SCSI) adapter is connected to a misconfigured SAS domain. Most likely a shelf in-port is attached to another in-port, or a shelf out-port is attached to another out-port. The SAS domain is effectively inoperative and attached SAS adapters will be taken offline. When the misconfiguration is corrected, the affected adapters can be brought back online.

#### Corrective Action

Correct the inter-shelf cabling configuration as defined in the applicable hardware documentation, and then bring the affected adapters back online.

## Syslog Message

Adapter "%s" is connected to misconfigured SAS domain and being taken offline.

## Parameters

**adapterName** (STRING): The name of the associated SAS host bus adapter.

## sas.config.badtransition

### Severity

ERROR

### Description

This message occurs when a SAS (Serial Attached SCSI) adapter is connected to a misconfigured SAS domain. A shelf with a 12 Gigabit-per-second (12 Gb/s) data rate might have been attached to a 3 or 6 Gb/s homogeneous stack or vice-versa. An external homogeneous stack does not allow for any transitions between 12 Gb/s and other data rates. The only exception is when an adapter is attached to an embedded internal shelf. A single transition between an internal shelf and the first external shelf attached to it is allowed. The message will appear the first time a transition is detected on a reboot or storage discovery operation.

### Corrective Action

Correct the inter-shelf cabling configuration as defined in the applicable hardware documentation. Ensure that the shelf stack is homogeneous. 3 Gb/s and 6 Gb/s shelves are considered homogeneous.

## Syslog Message

Unsupported speed transition by expander WWN: %s at level %d in SAS stack on adapter port %s.

## Parameters

**expanderWWN** (STRING): WWN of the expander that caused the illegal transition.

**expanderLevel** (INT): Level of the expander in the stack causing the transition. It is the physical depth of the shelf in the SAS stack where the error occurs, counting outwards from the adapter port.

**adapterName** (STRING): Name of the associated SAS host bus adapter port.

## sas.config.invalid

### Severity

ERROR

### Description

This message occurs when the SAS (Serial Attached SCSI) adapter is connected to a SAS domain that contains more than the supported number of disk shelves or enclosures. The SAS domain is functional, but some disk shelves/enclosures might not be visible to the system.

### Corrective Action

Remove disk shelves or enclosures attached to the applicable SAS adapter in question, until the supported limit is reached. (4 for FAS2040 and FAS2050 systems, 10 for all other systems)

## Syslog Message

Adapter "%s" is connected to more than the supported number of disk shelves or enclosures.

## Parameters

**adapterName** (STRING): The name of the associated SAS adapter.

## **sas.config.mixed.detected**

### Severity

ERROR

### Description

This message occurs when a SAS (Serial Attached SCSI) disk shelf contains a mixture of SAS drives, SATA drives or Bridged SAS drives. Mixing drive types within a disk shelf is not supported.

### Corrective Action

Ensure that each SAS disk shelf is populated with drives of only one type.

### Syslog Message

SAS shelf "%s" attached to adapter "%s" contains a mixture of drive types. Mixed configurations are not supported.

## Parameters

**shelfID** (STRING): ID or serial number of the disk shelf where a mixture of drive types was detected.

**adapterName** (STRING): Name of the associated SAS host bus adapter.

## **sas.cpr events**

### **sas.cpr.failed**

### Severity

EMERGENCY

### Description

This message occurs when the CPR (Core PCIe Resilience) infrastructure determines that the SAS controller hardware cannot be recovered, and the associated SAS ports will be taken offline.

### Corrective Action

Replace the SAS controller.

### Syslog Message

CPR recovery failed on SAS HBA %s.

## Parameters

**adapterName** (STRING): Name of the associated SAS host bus adapter (HBA).

### **sas.cpr.recoveryThreshold**

### Severity

EMERGENCY

### Description

This message occurs when the SAS driver determines that CPR (Core PCIe Resilience) recovery of the

SAS controller hardware has exceeded thresholds, and the associated SAS ports will be taken offline.

#### **Corrective Action**

Replace the SAS controller.

#### **Syslog Message**

CPR recovery threshold has been exceeded on SAS HBA %s.

#### **Parameters**

**adapterName** (STRING): Name of the associated SAS host bus adapter (HBA).

### **sas.cpr.recoveryTimeout**

#### **Severity**

ALERT

#### **Description**

This message occurs when the SAS driver determines that CPR (Core PCIe Resilience) recovery of the SAS controller hardware has timed out, and the associated SAS ports will be taken offline.

#### **Corrective Action**

Replace the SAS controller.

#### **Syslog Message**

CPR recovery timed out on SAS HBA %s.

#### **Parameters**

**adapterName** (STRING): Name of the associated SAS host bus adapter (HBA).

## **sas.device events**

### **sas.device.bad.vmdisk.header**

#### **Severity**

ERROR

#### **Description**

This message occurs when a disk's vmdisk header is bad. The disk is unusable by the system.

#### **Corrective Action**

Remove the disk from the system and either write a new vmdisk header to the disk or zero the disk.

#### **Syslog Message**

SAS adapter %s reports the device at port identifier %d has an invalid vmdisk header: %s.

#### **Parameters**

**adapterName** (STRING): Name of the associated SAS host bus adapter.

**devicePortId** (INT): Identifier of the SAS port device.

**errorString** (STRING): Identifies what is wrong with the vmdisk header.

## **sas.device.invalid.wwn**

### **Severity**

ERROR

### **Description**

The SAS device responded with an invalid World-Wide Name.

### **Corrective Action**

Power-cycling the device may allow it to recover from this problem.

### **Syslog Message**

SAS adapter %s reports the device at port identifier 0x%x has an invalid World-Wide Name.

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**devicePortId** (INTHEX): The SAS port identifier of the device.

## **sas.device.timeout**

### **Severity**

ERROR

### **Description**

All outstanding commands to the specified device have not completed within the allotted time. As part of the standard error handling sequence managed by the Data ONTAP SAS driver, all commands to the device will be aborted and reissued.

### **Corrective Action**

Device level timeouts are a common indication of a SAS link stability problem. In some cases the link is operating normally and the specified device is having trouble processing I/O requests in a timely manner. In such cases the specified device should be evaluated for possible replacement. Quite often the problem results from the partial failure of a component involved in the SAS transport. Common things to check include complete seating of drive carriers in enclosure bays, properly secured cable connections, IOM seating, and crimped or otherwise damaged cables.

### **Syslog Message**

Adapter %s encountered a device timeout on %s device %s.

### **Parameters**

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**deviceType** (STRING): The device type.

**deviceName** (STRING): The name of the SAS device.

## **sas.initialization events**

### **sas.initialization.failed**

### **Severity**

ERROR

## Description

The SAS adapter failed to initialize the link and appears to be unattached/disconnected.

## Corrective Action

If the adapter is in use check cabling and if connected to disk shelves check the seating of IOM cards.

## Syslog Message

Initialization failed on SAS adapter %s.

## Parameters

**adapterName** (STRING): The name of the associated SAS host bus adapter.

# sas.link events

## sas.link.error

### Severity

ERROR

### Description

The SAS adapter was unable to recover the link and is going offline.

### Corrective Action

If the adapter is in use check cabling and if connected to disk shelves check the seating of IOM cards and disks. If this does not resolve the issue please contact NetApp technical support for assistance in diagnosing a SAS issue.

### Syslog Message

Could not recover link on SAS adapter %s after %d seconds. Offlining the adapter.

### Parameters

**adapterName** (STRING): The name of the associated SAS host bus adapter.

**downTime** (INT): The amount of time that the link was down.

# sas.reset events

## sas.reset.exp.discovery.resp

### Severity

NOTICE

### Description

This message occurs when SAS adapter driver resets the expander because of failure status of discovery response.

### Corrective Action

If this message is infrequent, no action is needed. Otherwise, contact NetApp technical support.



## Syslog Message

SAS adapter [%s]: Incorrect status in discovery response, resetting expander: %s

## Parameters

**adapterName** (STRING): Name of the associated SAS adapter.

**expanderName** (STRING): Name of the expander being reset.

## **sas.reset.exp.open.dest**

### Severity

NOTICE

### Description

This message occurs when SAS adapter driver resets the expander because open frames are receiving excessive no/bad destination status.

### Corrective Action

If this message is infrequent, no action is needed. Otherwise, contact NetApp technical support.

## Syslog Message

SAS adapter [%s]: Excessive open no/bad destination status, resetting expander: %s

## Parameters

**adapterName** (STRING): Name of the associated SAS adapter.

**expanderName** (STRING): Name of the expander being reset.

## **sas.reset.exp.open.to**

### Severity

NOTICE

### Description

This message occurs when SAS adapter driver resets the expander because of excessive open timeouts.

### Corrective Action

If this message is infrequent, no action is needed. Otherwise, contact NetApp technical support.

## Syslog Message

SAS adapter [%s]: Excessive open timeouts occurred, resetting expander: %s

## Parameters

**adapterName** (STRING): Name of the associated SAS adapter.

**expanderName** (STRING): Name of the expander being reset.

## **sas.reset.exp.smp**

### Severity

NOTICE

## Description

This message occurs when SAS adapter driver resets the expander because of excessive SMP failures.

## Corrective Action

If this message is infrequent, no action is needed. Otherwise, contact NetApp technical support.

## Syslog Message

SAS adapter [%s]: Excessive SMP failures occurred, resetting expander: %s

## Parameters

**adapterName** (STRING): Name of the associated SAS adapter.

**expanderName** (STRING): Name of the expander being reset.

# sas.shelf events

## sas.shelf.conflict

### Severity

ERROR

### Description

This message occurs when the system detects that two or more SAS (serial attached SCSI) disk shelves have the same shelf ID. The SAS domain is functional, but references to disk shelves will be based on disk shelf serial numbers, not disk shelf IDs.

### Corrective Action

Reassign disk shelf IDs so that no conflict exists.

### Syslog Message

At least two SAS disk shelves have the same disk shelf ID.

### Parameters

(None).

# sas.unsupported events

## sas.unsupported.port

### Severity

ALERT

### Description

This message occurs when ONTAP detects that a cable is attached to an I/O module port that is permanently disabled and does not support attachment.

### Corrective Action

Remove the cable from the unsupported port on the I/O module.

## **Syslog Message**

Port 0b2 on shelf ID %lld.%s, cable detected on unsupported port.

## **Parameters**

**shelfID** (LONGINT): Shelf ID of the shelf reporting the error.

**element** (STRING): I/O module designator (A or B).

# sasmon events

## sasmon.adapter events

### sasmon.adapter.phy.disable

#### Severity

ERROR

#### Description

This message occurs when a Serial Attached SCSI (SAS) transceiver (phy) attached to a SAS host bus adapter (HBA) is disabled due to one of the following reasons: exceeded loss of double word synchronization error threshold; exceeded running disparity error threshold; exceeded invalid double word error threshold; exceeded phy reset problem threshold; exceeded broadcast change threshold.

#### Corrective Action

If the HBA is in use, check the cabling. If it is connected to the disk shelves, check the seating of the IOM (I/O module) cards. If that does not fix the problem, contact NetApp technical support.

#### Syslog Message

Phy %d of HBA %s will be disabled due to %s errors.

#### Parameters

**phyNum** (INT): Number of the phy being disabled.

**adapterName** (STRING): Name of the associated SAS HBA.

**errorType** (STRING): Type of error based on which the phy is being disabled.

## sasmon.disable events

### sasmon.disable.module

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Data ONTAP® module responsible for monitoring Serial Attached SCSI (SAS) domain's transient errors is disabled due to the environment variable `disable-sasmon?` being set to "true".

#### Corrective Action

Set the environment variable `disable-sasmon?` to "false" to enable this monitor module.

#### Syslog Message

SAS domain is not monitoring transient errors.

#### Parameters

(None).

# savecore events

## savecore.completed events

### savecore.completed

**Severity**

INFORMATIONAL

**Description**

A synchronous savecore operation has completed.

**Corrective Action**

(None).

**Syslog Message**

Savecore completed

**Parameters**

(None).

## savecore.nospace events

### savecore.nospace

**Severity**

ERROR

**Description**

Savecore has determined that the space available for storing the unsaved core is not sufficient.

**Corrective Action**

Delete old or unneeded files in the root filesystem. Check in `/etc/crash` (`/crash` on Netcaches) for unneeded cores. Make sure to delete any snapshots that include the deleted files.

**Syslog Message**

The available disk space (%u blocks) is not sufficient to save the core, which may need up to %d blocks.

**Parameters**

**avail** (INT): Number of blocks available for saving the core  
**needed** (INT): Number of blocks needed for saving the core

## savecore.readonly events

### savecore.readonly

**Severity**

NOTICE

**Description**

An attempt to run the savecore command was made during a read-only boot. This is not permitted.

**Corrective Action**

(None).

**Syslog Message**

Core dumps not saved in last resort read-only mode

**Parameters**

(None).

## savecore.waiting events

### savecore.waiting

**Severity**

INFORMATIONAL

**Description**

Savecore is running synchronously, and the console will be unresponsive until it is complete. This happens when savecore is given the "-f" or "-w" argument, or when previous attempts to save the core asynchronously fail.

**Corrective Action**

(None).

**Syslog Message**

Pausing to wait for savecore admin to finish

**Parameters**

(None).

# sblade events

## sblade.duplicate events

### sblade.duplicate.iscsi.addr

#### Severity

ALERT

#### Description

This message occurs when a duplicate IP address is detected on more than one iSCSI interface. Only one of the interfaces is sent as an iSCSI target portal.

#### Corrective Action

Delete all but one of the iSCSI interfaces which share an IP address.

#### Syslog Message

Vserver '%s' iSCSI interface '%s' has the same IP address '%s' as another iSCSI interface. This interface will not be sent as an iSCSI target portal.

#### Parameters

**vserver** (STRING): Vserver containing the iSCSI interfaces.

**interface** (STRING): Interface that contains a duplicate IP address.

**ipAddress** (STRING): IP address that was duplicated.

# scsi events

## scsi.cmd events

### scsi.cmd.aborted

#### Deprecated

This event has been removed as it does not provide any useful information for a customer.

#### Severity

ERROR

#### Description

The aborted event is issued if an I/O request is aborted by a class driver.

#### Corrective Action

(None).

#### Syslog Message

%s device %s: Command aborted: cdb %s (%d).

#### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command was aborted.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ(10) or a WRITE(10), the logical block address and the transfer length are included as well.

**dTime** (INT): The delta time from when the command was first issued.

### scsi.cmd.abortedByHost

#### Severity

ERROR

#### Description

This message occurs when the host adapter aborts an I/O request, either because it was a victim or because of an HA timeout.

#### Corrective Action

This condition does not necessarily mean that the target device is problematic. High workloads can cause loop saturation, leading to device contention for the bus. Transport issues can also cause loop throughput to decrease, thereby causing I/Os to take longer than normal. Seen occasionally, no action is required; the system handles the condition automatically by retrying the I/O that was aborted. Seen in association with a persistent error condition, such as device resets or transport errors, it is likely the symptom of a loop stability problem.

#### Syslog Message

%s device %s: Command aborted by host adapter: HA status 0x%x: cdb %s. %s

#### Parameters

**deviceType** (STRING): Device type.



**deviceName** (STRING): Electrical path name of the device on which the command failed.

**ha\_status** (INHEX): Value returned by the host adapter driver describing the completion status of the I/O request execution at the transport or physical link layer.

**cdb** (STRING): String containing the opcode byte of the CDB that was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.adapterFatal

### Severity

ERROR

### Description

The adapter fatal event is issued when the host adapter has detected a problem with either the buffer list or has exceeded the DMA size.

### Corrective Action

(None).

### Syslog Message

%s device %s: Adapter fatal error: HA status 0x%x: cdb %s.

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**haStatus** (INHEX): The value of the HA status field for the failed command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

## scsi.cmd.adapterHardwareError

### Severity

ERROR

### Description

The adapter hardware error is issued when the host adapter has detected a loop is down, the scsi bus transport had an error or a problem with the adapter hardware itself.

### Corrective Action

Evaluate for partial component failure the following loop hardware involved in the Fibre Channel transport: cables, transceivers, and I/O modules (LRC, ESH/ESH2, AT-FC, AT-FC2, and AT-FCX). If this error occurs in conjunction with transport errors on the same loop, ensure that all drive carriers are seated firmly in their enclosure bays, and verify that all cables are properly connected and undamaged. Also check the backplane connectors—if they are bent or broken, you must replace them. If the error persists for a specific device or I/O module from multiple adapters, you should replace that hardware module.

### Syslog Message

%s device %s: Adapter detected hardware error: HA status 0x%x: cdb %s.

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**haStatus** (INTHEX): The value of the HA status field for the failed command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

## scsi.cmd.adapterHardwareErrorEMSOOnly

### Severity

ERROR

### Description

This message occurs when the host adapter detects that a loop is down, the scsi bus transport has an error, or the adapter hardware itself has a problem, such as with a disk drive, cable, and so on.

### Corrective Action

Evaluate for partial component failure the following loop hardware involved in the Fibre Channel transport: cables, transceivers, and I/O modules (LRC, ESH/ESH2, AT-FC, AT-FC2, and AT-FCX). If this error occurs in conjunction with transport errors on the same loop, ensure that all drive carriers are seated firmly in their enclosure bays, and verify that all cables are properly connected and undamaged. Also check the backplane connectors—if they are bent or broken, you must replace them. If the error persists for a specific device or I/O module from multiple adapters, you should replace that hardware module.

### Syslog Message

%s device %s: Adapter detected hardware error: HA status 0x%x: cdb %s. %s

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): Physical path name of the device on which the command failed. It is in the format "path\_id.device\_id", where "path\_id := switch:port | adapter". For example, switch1:5:40 or 8b.40.

**haStatus** (INTHEX): Value of the HA status field for the failed command. For this error message, the value is 6.

**cdb** (STRING): String containing the opcode byte of the command descriptor block (CDB) that was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are also included. For example, a hardware error during READ is in the format "0x28:logical\_block\_address:transfer\_length".

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.adapterResources

### Severity

ERROR

### Description

The adapter resource event is issued when the host adapter has run out of resources.

### Corrective Action

(None).

## Syslog Message

%s device %s: Adapter detected resource issue: HA status 0x%x: cdb %s.

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**haStatus** (INTHEX): The value of the HA status field for the failed command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

## scsi.cmd.checkCondition

### Severity

ERROR

### Description

This message occurs when a check condition is received from a device. A check condition is the mechanism whereby a target device reports an informational condition or error status to the requesting host. Either the condition that generated this event is an error that occurred during execution of the command and was not cleared by retrying the request or an informational condition reporting status of the present operation or media state.

### Corrective Action

A target status of check condition normally indicates an error at the device during execution of the requested command. Such cases are often the result of an intermittent device hardware or firmware problem that is automatically handled by the Data ONTAP drivers through command retries. In cases of repeated events, the specified device should be evaluated for proper operation and possible repair or replacement. A check condition with a SenseKey of "no sense" indicates an informational condition that is automatically handled by the Data ONTAP driver associated with the target device. These are normally not reported as error events. However, on occasion an unexpected informational condition may be reported by a target device. These cases should not be interpreted as a failure of the target device.

## Syslog Message

%s device %s: Check Condition: CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d). %s

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.checkCondition.ATATimeout

### Severity

ERROR

### Description

An ATATimeout error typically indicates a problem in communicating with a disk in an ATFCX-bridged shelf.

### Corrective Action

An initial ATA I/O Timeout will result in a SOFT RESET being issued to the disk in question. Subsequent ATA I/O Timeouts will result in a HARD RESET being performed to the disk followed by a POWER CYCLE of the disk. When any of these errors are encountered, ATFCX shelflog information is also updated. Check shelflog dynamic and persistent log information for controller details on the event. Check environment shelf\_stats output for additional error statistics from the disk in question. Check disk shm\_stats and disk shm\_stats ha.disk\_id for SMART log information from the disk in question. Collating this information will help determine if this is an isolated disk problem or a controller related problem. If the environment shelf\_stats are clean along with the shelflog entries, then the error is probably a misbehaving disk. If the problem continues, RMA of the disk is recommended. However, if shelf\_stats are not clean and/or shelflog errors are present, RMA of the ATFCX module is recommended.

### Syslog Message

%s device %s: Check Condition: ATA I/O Timeout CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

## scsi.cmd.checkCondition.ATATimeout.Failure

### Severity

ERROR

### Description

An ATATimeout error typically indicates a problem in communicating with a disk in an ATFCX-bridged shelf. A Failure message indicates that error recovery was unable to recover access to the drive and hence it is being actively failed.

### Corrective Action

An initial ATA I/O Timeout will result in a SOFT RESET being issued to the disk in question. Subsequent ATA I/O Timeouts will result in a HARD RESET being performed to the disk followed by a POWER CYCLE of the disk. When any of these errors are encountered, ATFCX shelflog information is also updated. Check

shelflog dynamic and persistent log information for controller details on the event. Check environment shelf\_stats output for additional error statistics from the disk in question. Check disk shm\_stats and disk shm\_stats ha.disk\_id for SMART log information from the disk in question. Collating this information will help determine if this is an isolated disk problem or a controller related problem. If the environment shelf\_stats are clean along with the shelflog entries, then the error is probably a misbehaving disk. If the problem continues, RMA of the disk is recommended. However, if shelf\_stats are not clean and/or shelflog errors are present, RMA of the ATFCX module is recommended.

### Syslog Message

%s device %s: Check Condition: ATA I/O Timeout - Failure CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

## scsi.cmd.checkCondition.ATATimeout.HardReset

### Severity

ERROR

### Description

An ATATimeout error typically indicates a problem in communicating with a disk in an ATFCX-bridged shelf. A Hard Reset will be performed to the drive to attempt to clear the problem.

### Corrective Action

An initial ATA I/O Timeout will result in a SOFT RESET being issued to the disk in question. Subsequent ATA I/O Timeouts will result in a HARD RESET being performed to the disk followed by a POWER CYCLE of the disk. When any of these errors are encountered, ATFCX shelflog information is also updated. Check shelflog dynamic and persistent log information for controller details on the event. Check environment shelf\_stats output for additional error statistics from the disk in question. Check disk shm\_stats and disk shm\_stats ha.disk\_id for SMART log information from the disk in question. Collating this information will help determine if this is an isolated disk problem or a controller related problem. If the environment shelf\_stats are clean along with the shelflog entries, then the error is probably a misbehaving disk. If the problem continues, RMA of the disk is recommended. However, if shelf\_stats are not clean and/or shelflog errors are present, RMA of the ATFCX module is recommended.

### Syslog Message

%s device %s: Check Condition: ATA I/O Timeout - Hard Reset CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**IFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

## scsi.cmd.checkCondition.ATATimeout.InternalReset

### Severity

ERROR

### Description

An ATATimeout error typically indicates a problem in communicating with a disk in an ATFCX-bridged shelf. An internal reset (aka power cycle) will be performed to the drive to attempt to clear the problem.

### Corrective Action

An initial ATA I/O Timeout will result in a SOFT RESET being issued to the disk in question. Subsequent ATA I/O Timeouts will result in a HARD RESET being performed to the disk followed by a POWER CYCLE of the disk. When any of these errors are encountered, ATFCX shelflog information is also updated. Check shelflog dynamic and persistent log information for controller details on the event. Check environment shelf\_stats output for additional error statistics from the disk in question. Check disk shm\_stats and disk shm\_stats ha.disk\_id for SMART log information from the disk in question. Collating this information will help determine if this is an isolated disk problem or a controller related problem. If the environment shelf\_stats are clean along with the shelflog entries, then the error is probably a misbehaving disk. If the problem continues, RMA of the disk is recommended. However, if shelf\_stats are not clean and/or shelflog errors are present, RMA of the ATFCX module is recommended.

### Syslog Message

%s device %s: Check Condition: ATA I/O Timeout - Internal Reset CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**IFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

## **scsi.cmd.checkCondition.ATATimeout.SoftReset**

### **Severity**

ERROR

### **Description**

An ATATimeout error typically indicates a problem in communicating with a disk in an ATFCX-bridged shelf. A Soft Reset will be performed to the drive to attempt to clear the problem.

### **Corrective Action**

An initial ATA I/O Timeout will result in a SOFT RESET being issued to the disk in question. Subsequent ATA I/O Timeouts will result in a HARD RESET being performed to the disk followed by a POWER CYCLE of the disk. When any of these errors are encountered, ATFCX shelflog information is also updated. Check shelflog dynamic and persistent log information for controller details on the event. Check environment shelf\_stats output for additional error statistics from the disk in question. Check disk shm\_stats and disk shm\_stats ha.disk\_id for SMART log information from the disk in question. Collating this information will help determine if this is an isolated disk problem or a controller-related problem. If the environment shelf\_stats are clean along with the shelflog entries, then the error is probably a misbehaving disk. If the problem continues, RMA of the disk is recommended. However, if shelf\_stats are not clean and/or shelflog errors are present, RMA of the ATFCX module is recommended.

### **Syslog Message**

%s device %s: Check Condition: ATA I/O Timeout - Soft Reset CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

### **Parameters**

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): The integer value of the fru for the failed command.

**DTime** (INT): Delta time from first issued to sense logged.

## **scsi.cmd.checkCondition.SATA.DriveDown**

### **Severity**

ERROR

### **Description**

An SATA.DriveDown error typically indicates a problem in communicating with a disk behind a SAS to SATA bridge. A power cycle will be performed on the drive to attempt to clear the problem.

## Corrective Action

A SATA Drive Down condition should clear after a drive power cycle. If the problem persists, the drive will be failed at which point RMA of the disk is recommended.

## Syslog Message

device %s: Check Condition: SATA Drive Down CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

## Parameters

**deviceName** (STRING): The electrical path name of the device on which the command failed.  
**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.  
**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.  
**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.  
**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.  
**iASC** (INTHEX): The integer value of the sense code for the failed command.  
**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.  
**iFRU** (INTHEX): The integer value of the fru for the failed command.  
**DTime** (INT): Delta time from first issued to sense logged.

## scsi.cmd.checkCondition.SATA.Timeout

### Severity

ERROR

### Description

An SATA.Timeout error typically indicates a problem in communicating with a disk behind a SAS to SATA bridge. A hard reset will be performed on the drive to attempt to clear the problem.

### Corrective Action

A SATA Timeout condition should clear after a drive reset. If the problem persists, the drive report a drive down condition at which point it will be power cycle.

### Syslog Message

device %s: Check Condition: SATA Drive Timeout CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d).

### Parameters

**deviceName** (STRING): The electrical path name of the device on which the command failed.  
**cdb** (STRING): A string containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.  
**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.  
**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.  
**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.  
**iASC** (INTHEX): The integer value of the sense code for the failed command.  
**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.  
**iFRU** (INTHEX): The integer value of the fru for the failed command.  
**DTime** (INT): Delta time from first issued to sense logged.



## scsi.cmd.contingentAllegiance

### Severity

ERROR

### Description

This message occurs when an I/O request receives a check condition, but there is no sense data.

### Corrective Action

If there are command check conditions with no associated sense data with the device specified in this event and there are no other devices on the same bus/loop experiencing similar problems, the device should be evaluated for proper operation and possible repair/replacement.

### Syslog Message

%s device %s: Contingent allegiance: cdb %s. %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command received the contingent allegiance condition.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.driveFailureError

### Severity

ERROR

### Description

The drive reports failure after startup. Some i/o operation has failed on it and it should be removed.

### Corrective Action

(None).

### Syslog Message

%s device %s has failed during operation. Device %s is now unusable as a data or parity disk. It should be removed.

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The pathname to the device that has failed the operation.

**slotName** (STRING): The name of the slot in which the event occurred.

## scsi.cmd.excessiveVictim

### Severity

ERROR

## Description

This message occurs when the Storage Area Network (SAN) target or fabric continually aborts I/O operations. An initial victim abort results in the retry of the I/O operation without logging any messages. Subsequent victim aborts result in the logging of messages and additional retries. These messages are logged if there are continual aborts from the fabric or the target.

## Corrective Action

Check the fabric ports and the target ports for problems in any of the components, and if necessary, take corrective actions. Check the fabric and the target logs for any event that caused the fabric or the target to abort the I/O operations. If the problem persists, troubleshoot the fabric or target.

## Syslog Message

device %s: excessive victim abort: delta time %d: retry CDB %s: victim retry count %d: retry count %d: timeout retry count %d: path retry count %d: adapter status 0x%0x: target status 0x%0x: sense data %s - %s (0x%x - 0x%x 0x%x ). %s

## Parameters

**deviceName** (STRING): Path name of the device on which the command failed.  
**DTime** (INT): Time, in milliseconds, from when the command was first issued until this message was logged.  
**cdb** (STRING): String containing the opcode byte of the Command Descriptor Block (CDB) that was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are also included.  
**victimRetryCount** (INT): Number of retries due to I/O failure not related to this I/O operation.  
**retryCount** (INT): Number of retries from the SCSI layer.  
**timeoutRetryCount** (INT): Number of times the I/O operation was timed out and retried.  
**pathRetryCount** (INT): Number of retries for the path.  
**adapterStatus** (INTHEX): Host adapter status.  
**targetStatus** (INTHEX): Target status.  
**sSenseKey** (STRING): String containing a description of the sense key for the failed command.  
**sSenseCode** (STRING): String containing a description of the sense code for the failed command. If no description is available, this is an empty string.  
**iSenseKey** (INTHEX): Integer value of the sense key for the failed command.  
**iASC** (INTHEX): Integer value of the sense code for the failed command.  
**iASCQ** (INTHEX): Integer value of the sense code qualifier for the failed command.  
**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.failDevice

### Severity

ERROR

### Description

The failed device event is issued when a device emits a vendor specific sense key.

### Corrective Action

The failed device should be replaced immediately.

### Syslog Message

- ALERT! \*\* %s device %s has failed! Replace device at once!

## Parameters

**deviceType** (STRING): The type of device which has failed.

**deviceName** (STRING): The electrical path name of the device which failed.

## scsi.cmd.floatTimeout

### Severity

NOTICE

### Description

This message occurs when an Azure page blob I/O operation fails due to a timeout. ONTAP retries the I/O command for a maximum of five retries, until the command succeeds or fails. Successive retries increase the timeout period.

### Corrective Action

(None).

### Syslog Message

"%s" device "%s": The page blob I/O operation timeout increased to "%d" seconds: HA status "0x%x": OSC Error "%d": retry count "%d": elapsed time "%d" milliseconds: cdb %s. %s

## Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**basicTimeout** (INT): Basic timeout, in seconds, set for the I/O operation.

**haStatus** (INHEX): Value of the high-availability (HA) status field for the failed command.

**oscError** (INT): Value of the Object Store Client (OSC) error for the failed command.

**retryCount** (INT): Number of retries from the SCSI layer.

**ETime** (INT): Time, in milliseconds, from when the command was first issued until this I/O operation failed.

**cdb** (STRING): String containing the opcode byte of the Command Descriptor Block (CDB) that was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**diskInformation** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.fw.download.failed.butReverted

### Severity

ERROR

### Description

This event indicates that the download process was unsuccessful and thus the original running firmware is, again, the current running version.

### Corrective Action

If this error continues to be encountered, replace the FC-AT module in the shelf.

### Syslog Message

Firmware download failed, but successfully reverted to the previous running version on %s%d

## Parameters

**shelfName** (STRING): The name of the shelf in which the firmware download status is being provided.

**shelf\_num** (INT): The value of the shelf number.

## scsi.cmd.fw.download.failed.notRecoverable

### Severity

ERROR

### Description

This event indicates that the download process failed and the previous running version of the firmware was not recoverable.

### Corrective Action

Replace the FC-AT module in the shelf.

### Syslog Message

Firmware download was unsuccessful and the previous running shelf firmware was not recoverable on %s%d. Once the shelf is power cycled, the FC-AT module may not be operational.

## Parameters

**shelfName** (STRING): The name of the shelf in which the firmware download status is being provided.

**shelf\_num** (INT): The value of the shelf number.

## scsi.cmd.fw.download.successful

### Severity

INFORMATIONAL

### Description

This event indicates that the firmware download process completed successfully.

### Corrective Action

(None).

### Syslog Message

Firmware download was successful on %s%d

## Parameters

**shelfName** (STRING): The name of the shelf in which the firmware download status is being provided.

**shelf\_num** (INT): The value of the shelf number.

## scsi.cmd.fw.download.successful.onRetry

### Severity

NOTICE

### Description

This event indicates that the download process was successful after being retried.

## Corrective Action

(None).

## Syslog Message

Firmware download was successful after being retried on %s%d

## Parameters

**shelfName** (STRING): The name of the shelf in which the firmware download status is being provided.

**shelf\_num** (INT): The value of the shelf number.

## scsi.cmd.noMorePaths

### Severity

ERROR

### Description

The no more paths event is issued when all paths to a device have been tried without success.

### Corrective Action

(None).

### Syslog Message

%s device %s: No more paths to device: cdb %s. All retries have failed.

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

## scsi.cmd.notReadyCondition

### Severity

NOTICE

### Description

This message occurs when an I/O request receives a check condition where the device reports that it is not ready to process the I/O request. The I/O request is delayed to wait for the device to become ready, and then it is tried again.

### Corrective Action

If command not ready check conditions associated with the device specified in this event persist, the device should be evaluated for proper operation and possible repair/replacement.

### Syslog Message

%s device %s: Device returns not yet ready: CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x)(%d). %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**sSenseKey** (STRING): A string containing a description of the sense key for the failed command.

**sSenseCode** (STRING): A string containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): The integer value of the sense key for the failed command.

**iASC** (INTHEX): The integer value of the sense code for the failed command.

**iASCQ** (INTHEX): The integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): The integer value of the fru for the failed command.

**dTime** (INT): The delta time from when the command was first issued.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.overrun

### Severity

ERROR

### Description

This message occurs when an I/O request receives an overrun error. Too much data was received. I/O will be retried.

### Corrective Action

If there are command overruns associated with the device specified in this event and there are no other devices on the same bus/loop experiencing similar problems, the device should be evaluated for proper operation and possible repair/replacement.

### Syslog Message

%s device %s: Received a data overrun: cdb %s. Too much data was received. Possible transmission error. I/O will be retried. %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.pastTimeToLive

### Severity

ERROR

### Description

This message occurs when an I/O request is not completed in the maximum time allowed. It is possible that the target device is attempting an error recovery operation that is impeding the progress of other I/O requests. Often these conditions are not permanent and can be cleared by retrying the request. However, the system must limit the maximum time it waits for a device to successfully complete a request. Similarly, the number of retries is also limited. If the number of retries is not exhausted before the maximum time

allowed is exceeded, this event is generated.

### Corrective Action

If there are command timeouts associated with the device specified in this event and there are no other devices on the same bus/loop experiencing similar problems, the device should be evaluated for proper operation and possible repair/replacement.

### Syslog Message

%s device %s: request failed after try #%d: cdb %s. %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**retryCount** (INT): The number of times the command was retried before its timer expired.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.requestSenseFailed

### Severity

ERROR

### Description

The request sense failed event is issued when the HA can not the auto request sense for a command which has failed.

### Corrective Action

(None).

### Syslog Message

%s device %s: Request sense failed: out status 0x%x: cdb %s.

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cmdStatus** (INTHEX): The error flags which were set for the command which failed.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

## scsi.cmd.selectionTimeout

### Severity

ERROR

### Description

This message occurs when the host adapter returns either a selection time out error (SCSI) or an invalid target error (FC). In effect, the specified device failed to acknowledge a request.

## Corrective Action

Either the targeted device failed to respond or the address of the targeted device is no longer valid. If an alternate path to the device exists, the I/O will be automatically retried on the alternate path. If seen in conjunction with other timeouts on the same bus/loop, there may be a problem with the bus/loop, not the specific device. If this event is not seen in conjunction with any other error recovery activity on the same bus/loop, the specified device should be evaluated for proper operation and possible replacement.

## Syslog Message

%s device %s: Adapter/target error: HA status 0x%x: cdb %s. Targeted device did not respond to requested I/O. I/O will be retried. %s

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**haStatus** (INHEX): The value of the HA status field for the failed command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.senseDataDump

### Severity

ERROR

### Description

This message occurs when a device emits a vendor-specific sense key. The returned sense data is dumped to provide the vendor with enough information to track down the problem.

### Corrective Action

Because the sense key is vendor-specific, the vendor of the device should be contacted for any corrective actions.

### Syslog Message

class 0x%x segment 0x%x sense\_key 0x%x info 0x%x 0x%x 0x%x 0x%x length 0x%x resv\_1 0x%x sense\_code 0x%x sense\_code\_qualifier 0x%x fru\_failed 0x%x flags 0x%x field\_pointer 0x%x. %s

### Parameters

**class** (INHEX): The response code byte of the request sense data.

**segment** (INHEX): The segment number byte of the request sense data.

**senseKey** (INHEX): The Sense key byte of the request sense data.

**info0** (INHEX): The most significant byte of the information field of the request sense data.

**info1** (INHEX): The second most significant byte of the information field of the request sense data.

**info2** (INHEX): The third most significant byte of the information field of the request sense data.

**info3** (INHEX): The least significant byte of the information field of the request sense data.

**length** (INHEX): The additional sense length field of the request sense data.

**resv\_1** (INHEX): The command specific information field of the request sense data.

**senseCode** (INHEX): The additional sense code field of the request sense data.

**senseCodeQualifier** (INHEX): The additional sense code qualifier field of the request sense data.

**fruFailed** (INHEX): The field replaceable unit code field of the request sense data.

**flags** (INHEX): The flags byte of the sense key specific field of the request sense data.



**fieldPointer** (INTHEX): The field pointer bytes of the sense key specific field of the request sense data.  
**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.SFRPCrestartError

### Severity

EMERGENCY

### Description

This message occurs when an SFRPC I/O request failure is being retried and the request exceeds the maximum time allowed for the request ("Max Time to Live"). The system will restart to attempt recovery.

### Corrective Action

No additional action is required.

### Syslog Message

%s device %s: Transport error exceeded "Max Time to Live" limit during the execution of the command. The system will attempt to restart to clear the requests: HA status 0x%x: ffm\_status: %s:%d CDB %s (%d). %s

### Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**ha\_status** (INTHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**ffm\_statusName** (STRING): Name of the status returned by the FFM services describing the completion status of the I/O request at the FFM layer.

**ffm\_status** (INTHEX): Value returned by the FFM services describing the completion status of the I/O request at the FFM layer.

**cdb** (STRING): String containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a) command, the logical block address and the transfer length are also included.

**dTime** (INT): Delta time from when the command was first issued.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, serial number and slice ID.

## scsi.cmd.SFRPCtransportError

### Severity

ERROR

### Description

This message occurs when an SFRPC I/O request encounters a transport error during command execution. Transport errors are typically the result of the SFRPC layer's inability to complete a request successfully.

### Corrective Action

Transport errors are typically unexpected events. System software takes appropriate recovery actions by retrying the I/O. No additional action is required.

### Syslog Message

%s device %s: Transport error during execution of command: HA status 0x%x: ffm\_status: %s:%d CDB %s (%d). %s

## Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**ha\_status** (INHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**ffm\_statusName** (STRING): Name of status returned by the FFM services describing the completion status of the I/O request at the FFM layer.

**ffm\_status** (INHEX): Value returned by the FFM services describing the completion status of the I/O request at the FFM layer.

**cdb** (STRING): String containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a) command, the logical block address and the transfer length are included as well.

**dTime** (INT): Delta time from when the command was first issued.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, serial number and slice ID.

## scsi.cmd.slotfailureError

### Severity

ERROR

### Description

The slot failure error event is issued if an I/O request got (0x02, 0x04, 0x05) returned from the sense information.

### Corrective Action

(None).

### Syslog Message

%s device %s has detected a slot failure. Device %s is now inaccessible. Shelf must be replaced to resolve problem.

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**slotName** (STRING): The name of the slot in which the event occurred.

## scsi.cmd.targetStatus

### Severity

ERROR

### Description

This message occurs when an I/O request receives a status that is neither a reservation conflict nor a check condition, and represents an error that can not be retried.

### Corrective Action

If the command target status associated with the device specified in this event persists and/or the device is failed, the device should be evaluated for proper operation and possible repair/replacement.

## Syslog Message

%s device %s: Target Status %s: cdb %s. %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**targetStatus** (STRING): A string containing the target status received from the command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.transportError

### Severity

ERROR

### Description

This message occurs when an I/O request encounters a transport error during command execution.

Transport errors are typically the result of an incomplete data transfer (e.g., device expected more data than was received from the host adapter, or the adapter expected more data than was received from the device), a data overrun meaning that more data was transferred than was expected, a physical link transmission (parity or CRC) error, or a bus protocol error (phase error or an unexpected bus state transition.)

### Corrective Action

Transport errors are typically unexpected events. The exception is when drives are physically added or removed from the bus or drive enclosure. However, transport protocols do account for detection and recovery of transmission errors through CRC or Parity mechanisms. When such an error is reported to the adapter driver, the associated request is retried. For example, on a Fibre Channel link, an occasional transmission error is consistent with the expected Bit Error Rate specifications of the link. Output from the `fcadmin link_stats` command can be used to obtain error counts for CRC and transport errors. If they exceed observed values from normally operating loops, loop hardware should be evaluated for partial failure of a component involved in the Fibre Channel transport. Common things to check include complete seating of drive carriers in enclosure bays, properly secured cable connections, LRC/ESH seating, and crimped or otherwise damaged cables.

## Syslog Message

%s device %s: Transport error during execution of command: HA status 0x%x: cdb %s.

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**ha\_status** (INHEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (x02a) command, the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.transportErrorEMSONly

### Severity

ERROR

### Description

This message occurs when an I/O request encounters a transport error during command execution. Transport errors are typically unexpected events, unless you added or removed disk drives from the bus or drive enclosure. Unexpected transport errors are typically the result of the following conditions: An incomplete data transfer where the device expected more data from the host adapter than it received (or the adapter expected more data from the device than it received); A data overrun where more data than expected was transferred; A physical link transmission (parity or CRC [cyclic redundancy check]) error; Or a bus protocol error (phase error or unexpected bus state transition).

### Corrective Action

Transport protocols do account for detection and recovery of transmission errors through CRC or parity mechanisms. When this error is reported to the adapter driver, the associated request is retried. For example, on a Fibre Channel link, an occasional transmission error is consistent with the expected Bit Error Rate specifications of the link. Therefore, use output from the "fcadmin link\_stats" command to obtain error counts for CRC and transport errors. If the error counts exceed observed values from normally operating loops, evaluate for partial component failure the loop hardware involved in the Fibre Channel transport. Common things to check include complete seating of drive carriers in enclosure bays, properly secured cable connections, complete LRC/ESH module seating, and crimped or otherwise damaged cables.

### Syslog Message

%s device %s: Transport error during execution of command: HA status 0x%x: cdb %s. %s

### Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): Physical path name of the device on which the command failed. It is in the format "path\_id.device\_id", where "path\_id := switch:port | adapter". For example, switch1:5:40 or 8b.40.

**ha\_status** (INHEX): Value returned by the host adapter driver describing completion status of the I/O request execution at the transport layer or physical link layer. For this error message, the value is 9.

**cdb** (STRING): String containing the opcode byte of the command descriptor block (CDB) that was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a) command, the logical block address and the transfer length are also included. For example, 0x28:logical\_block\_address:transfer\_length.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.underrun

### Severity

ERROR

### Description

This message occurs when an I/O request receives an underrun error and has not specifically instructed the SCSI layer to ignore underrun errors.

### Corrective Action

If there are command underruns associated with the device specified in this event and there are no other devices on the same bus/loop experiencing similar problems, the device should be evaluated for proper operation and possible repair/replacement.

## Syslog Message

%s device %s: Received a data underrun: cdb %s. Not all the data was received. Possible transmission error. I/O will be retried. %s

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.cmd.unexpectedAdapterError

### Severity

ERROR

### Description

This message occurs when a hardware adapter returns an error that is unknown to the SCSI layer.

### Corrective Action

The associated hardware adapter should be evaluated for proper operation and possible repair/replacement.

## Syslog Message

%s device %s: Unexpected adapter error: HA status 0x%x: cdb %s. %s

## Parameters

**deviceType** (STRING): The device type.

**deviceName** (STRING): The electrical path name of the device on which the command failed.

**haStatus** (INHEX): The value of the HA status field for the failed command.

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (0x2a), the logical block address and the transfer length are included as well.

**disk\_information** (STRING): Formatted information about the disk, including the disk's vendor, model, firmware revision, and serial number.

## scsi.mcc events

### scsi.mcc.adt.ioTransportError

### Severity

ERROR

### Description

Generated when an I/O request encounters an error because of an iSCSI Protocol Error, a end device error, or some stack internal error.

## Corrective Action

Transport errors are typically unexpected events; the exception being when drives are physically added or removed from the bus or drive enclosure. However, transport protocols do account for detection and recovery of transmission errors through CRC or Parity mechanisms. When such an error is reported to the adapter driver, the associated request is retried. Common things to check include complete seating of drive carriers in enclosure bays, properly secured cable connections, LRC/ESH seating, and crimped or otherwise damaged cables.

## Syslog Message

mcc\_adt[%u] - Transport error during execution of command: HA status 0x%x: CAM transport status 0x%x : cdb %s.

## Parameters

**path\_instance** (INTEX): Value returned by the host adapter driver describing path instance.

**ha\_status** (INTEX): Value returned by the host adapter driver describing completion status of I/O request execution at the transport or physical link layer.

**cam\_status** (INTEX): Value returned by the host adapter driver in the FreeBSD CAM transport layer

**cdb** (STRING): A string containing the opcode byte of the CDB which was sent to the device. If the opcode was for a READ (0x28) or a WRITE (x02a) command, the logical block address and the transfer length are included as well.

# scsi.path events

## scsi.path.excessiveErrors

### Severity

ERROR

### Description

This message occurs when excessive errors are encountered on this device within a short-enough time period to raise concern that there might be a faulty component between the storage initiator and device. You can use the option 'ra.path\_switch.threshold' to control how sensitive Data ONTAP® is to device errors. Setting it to a higher value requires more errors to trigger this event.

### Corrective Action

Device errors are a common indication of storage domain stability problems. In some cases, the interconnect is operating normally and the specified device is having trouble processing I/O requests in a timely manner. In such cases, evaluate the specified device for possible replacement. Quite often the problem results from the partial failure of a component involved in the storage transport. Common things to check include complete seating of drive carriers in enclosure bays, properly secured cable connections, shelf IOM seating, failing Fibre Channel switch components, failing FC SAS bridge components and crimped or otherwise damaged cables.

### Syslog Message

Excessive errors encountered by adapter %s on %s device %s.

### Parameters

**adapterName** (STRING): Name of the associated storage host bus adapter.

**deviceType** (STRING): Device type.

**deviceName** (STRING): Name of the storage device.

## scsi.scan events

### scsi.scan.bus.halted

#### Severity

NOTICE

#### Description

This message occurs when the creation of the SCSI transport (xpt) path fails.

#### Corrective Action

Check whether the device is connected correctly.

#### Syslog Message

The path that is creating the xpt path has failed with status 0x%x. Bus scan halted.

#### Parameters

**status** (INTEX): SCSI scan bus halt status.

### scsi.scan.lun.halted

#### Severity

NOTICE

#### Description

This message occurs when the creation of the SCSI transport (xpt) path fails.

#### Corrective Action

Check whether the device is connected correctly.

#### Syslog Message

The path that is creating the xpt path has failed with status 0x%x. Bus scan halted.

#### Parameters

**status** (INTEX): SCSI scan LUN halted status.

## scsi.security events

### scsi.security.authLockOut

#### Deprecated

This event is removed to disassociate from SCSI, with the advent of NVMe encrypting drives. See replacement message EMS\_nse\_authLockedOut.

#### Severity

ERROR

#### Description

This message occurs when a self-encrypting drive (SED) reports excessive sequential failed authentication attempts with an incorrect Authentication Key (AK). To prevent brute-force AK attacks, the SED will no

longer accept even the correct AK. Access and control of the SED are affected. BandMaster1 authority lockout affects protection controls on user data. The system cannot authenticate using the data AK. A device power-cycle will result in data loss on the drive. For other lockouts, some aspect of compliance with Federal Information Processing Standard (FIPS) PUB 140-2 is compromised. Erase Master lockout disables cryptographic sanitize and destroy.

### Corrective Action

For BandMaster1 authority lockout, reset the AK to its default value and simultaneously erase user data using the "storage encryption disk sanitize" command or the (advanced privilege) "disk encrypt sanitize" command in node shell/maintenance mode. Power-cycle clears Physical Ownership lockout in some SEDs. For other lockouts, reset the SED to its as-manufactured state and erase user data using the "storage encryption disk revert-to-original-state" command or the (advanced privilege) "disk encrypt revert\_original" command in node shell/maintenance mode with the PSID value from the device label. If the PSID is not available, replace the drive. After clearing lockout, make the drive a spare by using the "storage disk assign" command and (advanced privilege) "storage disk unfail [-spare]" command as needed. To prevent lockouts, ensure that all cluster NSE nodes have all SED AKs. Display SED key IDs by using the "storage encryption disk show [-fips]" command. Verify them against those retrieved from the AK storage method: KMIP servers: Use the "security key-manager restore" command. Onboard key management: Use the "security key-manager key show" command.

### Syslog Message

Lockout on Storage Encryption device %s; security provider: %s, authority: %s, during operation "%s".

### Parameters

**deviceName** (STRING): Name of the device.

**securityProvider** (STRING): Name of the security provider.

**authority** (STRING): Name of the authority.

**operation** (STRING): Description of the operation that attempted authentication.

## scsi.sim events

### scsi.sim.checkCondition

#### Severity

ERROR

#### Description

A Check Condition is the mechanism whereby a target device reports an informational condition or error status to the requesting host. Either the condition that generated this event is an error that occurred during execution of the command and was not cleared by retrying the request or an informational condition reporting status of the present operation or media state.

#### Corrective Action

A target status of Check Condition normally indicates an error at the device during execution of the requested command. Such cases are often the result of an intermittent device hardware or firmware problem that is automatically handled by the Data ONTAP drivers through command retries. In cases of repeated events, the specified device should be evaluated for proper operation and possible repair or replacement. A Check Condition with a SenseKey of "no sense" indicates an informational condition that is automatically handled by the Data ONTAP driver associated with the target device. These are normally not reported as error events. However, on occasion an unexpected informational condition may be reported by a target device. These cases should not be interpreted as a failure of the target device.



## Syslog Message

%s device %s: Check Condition: CDB %s: Sense Data %s - %s (0x%x - 0x%x 0x%x 0x%x) Retry: %d/%d (%d).

## Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**cdb** (STRING): String containing the opcode byte of the Command block which encountered the error. For opcodes of READ (0x28) or WRITE (0x2a), the logical block address and the transfer length are also included.

**sSenseKey** (STRING): String containing a description of the sense key for the failed command.

**sSenseCode** (STRING): String containing a description of the sense code for the failed command. If no description is available, this is a null string.

**iSenseKey** (INTHEX): Integer value of the sense key for the failed command.

**iASC** (INTHEX): Integer value of the sense code for the failed command.

**iASCQ** (INTHEX): Integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): Integer value of the fru for the failed command.

**retryCount** (INT): Number of times the command was retried

**freeRetryCount** (INT): Number of free retries the command was given before it was successful.

**DTime** (INT): Delta time from first issued to sense logged.

# scsiblade events

## scsiblade.autoheal events

### scsiblade.autoheal.disabled

**Severity**

NOTICE

**Description**

This message occurs when the node loses connectivity with the majority of the cluster and autoheal has been explicitly disabled by the user. When this occurs, both FCP and iSCSI traffic are disabled on this node and require manual intervention to come back into service.

**Corrective Action**

(None).

**Syslog Message**

This node lost connectivity with the majority of the cluster. The option "san.autoheal.enable" is "off", so the cluster's SAN configuration will not automatically recover. Both FCP and iSCSI traffic are disabled on this node.

**Parameters**

(None).

## scsiblade.cache events

### scsiblade.cache.commit.fail

**Severity**

ALERT

**Description**

This message occurs when the SAN cache encounters an error that might cause configuration inconsistencies within the cluster. SAN data traffic might be interrupted until the problem is addressed.

**Corrective Action**

Reload the scsi-blade by using the (privilege: diagnostic) "debug san vserver reload" command for all Vservers on this node, or perform a planned takeover/giveback of this node.

**Syslog Message**

The SAN cache might be inconsistent. SAN data traffic might be interrupted.

**Parameters**

(None).

## scsiblade.disabled events

## **scsiblade.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the local node is made ineligible. When this occurs, both FCP and iSCSI traffic are disabled on the local node.

### **Corrective Action**

Modify the node's eligibility by using the (privilege: advanced) "cluster modify -eligibility true" command.

### **Syslog Message**

The local node is ineligible. FCP and iSCSI access is disabled.

### **Parameters**

(None).

## **scsiblade.dpa events**

### **scsiblade.dpa.disabled**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the node loses connectivity with the majority of the cluster and the Data Plane Availability feature has been explicitly disabled by the user. When this occurs, both FCP and iSCSI traffic are disabled on this node until connectivity is restored.

#### **Corrective Action**

(None).

#### **Syslog Message**

This node lost connectivity with the majority of the cluster. The option "san.dpa.enable" is "off", so both FCP and iSCSI traffic are disabled on this node until connectivity is restored.

#### **Parameters**

(None).

## **scsiblade.ha events**

### **scsiblade.ha.partner.dblade.id.retrieve.error**

#### **Severity**

ERROR

#### **Description**

Unable to retrieve HA Partner DBlade ID.

### Corrective Action

No corrective action

### Syslog Message

Unable to retrieve HA Partner DBLade ID.

### Parameters

(None).

## scsiblade.illegal events

### scsiblade.illegal.xfer.len

#### Severity

ERROR

#### Description

This message occurs when an error is returned for a SCSI command that exceeds the LUN's maximum supported transfer length of 1MB.

#### Corrective Action

This error is due to a problem on the host. Use one of the following recovery actions to address the problem. 1) Rescan the LUNs or reboot the host to enable discovery of the supported maximum transfer length. 2) Disable the host application that is causing the problem. 3) Report the bug to the host operating system or application provider.

#### Syslog Message

Returned Illegal Request response to host type %s for command opcode 0x%x with a transfer length of %lld bytes on LUN %s.

#### Parameters

**host\_os\_type** (STRING): Host OS type.  
**opcode** (INT): SCSI opcode.  
**length** (LONGINT): Request transfer length.  
**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.import events

### scsiblade.import.attr.mod

#### Severity

ERROR

#### Description

This message occurs when attributes of a SAN LUN under LUN import are modified. This is not a supported operation and the LUN import session needs to be deleted and re-created.

#### Corrective Action

Run the "lun import delete" and "lun import create" commands to re-create the import session.

## Syslog Message

Attributes of the following SAN LUN under LUN import are modified: Vserver %s, volume with MSID %llu, LUN serial number %s, LUN UUID %s. This is not a supported operation and the LUN import session needs to be deleted and re-created.

## Parameters

**vserver** (STRING): Name or UUID of the Vserver.  
**volumeMSID** (LONGINT): MSID of the container volume.  
**LUN\_Serial** (STRING): LUN serial number.  
**LUN\_UUID** (STRING): LUN UUID.

## scsiblade.import.dblade.fail

### Severity

ERROR

### Description

This message occurs when a SAN LUN import failover event fails for a Vserver during takeover or giveback. Some LUN import sessions configured on the Vserver will be affected. The failed LUN import sessions might not progress until a planned takeover or giveback is performed.

### Corrective Action

Perform a planned takeover or giveback.

## Syslog Message

SAN LUN import failover event failed for Vserver %s. Some LUN import sessions configured on the Vserver will be affected and will be in a failed state.

## Parameters

**vserver** (STRING): Name or UUID of the Vserver.

## scsiblade.import.lun.offline

### Severity

ERROR

### Description

This message occurs when a LUN currently being imported is brought offline by Data ONTAP® because of an I/O failure on the foreign disk. The issue is either that data cannot be read from or written to the source (foreign LUN), or that there is an internal error.

### Corrective Action

Verify the connectivity to the foreign disk. After resolving any connectivity issues, delete the LUN import and create it again.

## Syslog Message

LUN %s in volume with MSID %llu on Vserver %s has been brought offline due to an I/O failure on the underlying foreign disk with serial number %s.

## Parameters

**lunSerial** (STRING): Serial number of the LUN.  
**volumeMSID** (LONGINT): MSID of the volume.

**vserver** (STRING): Name or UUID of the Vserver.  
**foreignDiskSINo** (STRING): Serial number of the foreign disk.

## **scsiblade.import.volume.fail**

### **Severity**

ERROR

### **Description**

This message occurs when a SAN LUN import failover event fails for a volume during takeover or giveback. All LUN import sessions configured in this volume are affected. The failed LUN import sessions might not progress until a planned takeover or giveback is performed.

### **Corrective Action**

Perform a planned takeover or giveback.

### **Syslog Message**

SAN LUN import failover event failed for Vserver %s, volume with MSID %llu. Some LUN import sessions configured on the volume will be affected and will be in a failed state.

### **Parameters**

**vserver** (STRING): Name of the Vserver.  
**volumeMSID** (LONGINT): MSID of the volume.

## **scsiblade.init events**

### **scsiblade.init.abort**

#### **Severity**

ALERT

#### **Description**

This message occurs when the specified Vserver initialization was skipped due to a user-initiated abort. Giveback was allowed to succeed—in a degraded state—to allow further recovery. LUNs in the Vserver might not be accessible or might not be accessible on all paths.

#### **Corrective Action**

Resolve the previously identified fault for which the abort was issued, and then restart SAN initialization on the node by using the "debug san node restart-vserver-initialization" command.

#### **Syslog Message**

Vserver "%s" was not initialized due to a user-initiated abort request. SAN service for this Vserver might be degraded or unavailable on this node.

#### **Parameters**

**vserver** (STRING): Name of the Vserver.

### **scsiblade.init.failure**

**Severity**

ALERT

**Description**

This message occurs when there is a problem initializing the SAN Kernel Agent for one or more Vservers. LUNs in the affected Vserver(s) might not be accessible or might not be accessible via all paths.

**Corrective Action**

Perform a planned takeover of the node with the failure followed by a giveback.

**Syslog Message**

There was a problem initializing the SAN Kernel Agent for one or more Vservers. Vserver: %s. Reason: %s.

**Parameters**

**vserver** (STRING): Name of the Vserver for which the SAN Kernel Agent failed to initialize.

**error** (STRING): Error encountered while initializing the SAN Kernel Agent for the Vserver.

## scsiblade.invalid events

### scsiblade.invalid.xfer.len

**Severity**

NOTICE

**Description**

This message occurs when a SCSI command requests a transfer length greater than the LUN's maximum supported transfer length of 1MB. ONTAP® software completes the command normally but the host has sent this request in violation of the SCSI protocol.

**Corrective Action**

Rescan the LUNs or reboot the host to enable discovery of the supported maximum transfer length for the LUN.

**Syslog Message**

Host type %s sent command opcode 0x%x with a transfer length of %lld bytes, which exceeds the maximum of 1MB on LUN %s.

**Parameters**

**host\_os\_type** (STRING): Host OS type.

**opcode** (INT): SCSI opcode.

**length** (LONGINT): Request transfer length.

**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.isolated events

### scsiblade.isolated

**Severity**

ERROR

## Description

This message occurs when the node loses connectivity with the majority of the cluster. When this occurs, both FCP and iSCSI traffic are isolated to local access only. Remote cluster traffic is impacted. After connectivity is restored, the node automatically recovers and allows remote FCP and iSCSI access.

## Corrective Action

Correct the issue that caused the node to lose connectivity with the majority of the cluster. After the issue is corrected, remote service is restored.

## Syslog Message

This node lost connectivity with the majority of the cluster. When this occurs, both FCP and iSCSI traffic are isolated to local access only. Remote cluster traffic will be impacted until connectivity is restored.

## Parameters

(None).

# scsiblade.kernel.events

## scsiblade.kernel.volume.limbo.group

### Severity

EMERGENCY

### Description

This message occurs when the kernel cluster group for the specified volume is in an indeterminate state.

### Corrective Action

Contact NetApp technical support. For further information about correcting the problem, see the knowledgebase article 2017368.

### Syslog Message

The kernel cluster group for volume MSID 0x%llx of Vserver %s is in an indeterminate state. This will impact SAN management and configuration. Additionally, LUNs might not be accessible on this volume from this node.

### Parameters

**volume** (LONGINTEX): Volume identifier.

**vserver** (STRING): UUID of the Vserver.

## scsiblade.kernel.vserver.limbo.group

### Severity

EMERGENCY

### Description

This message occurs when the kernel cluster group for the specified Vserver is in an indeterminate state.

### Corrective Action

Contact NetApp technical support. For further information about correcting the problem, see the Knowledgebase article 2017367.



## Syslog Message

The kernel cluster group for Vserver %s is in an indeterminate state. This will impact SAN management and configuration. Additionally, LUNs might not be accessible on this Vserver from this node.

## Parameters

**vserver** (STRING): UUID of the Vserver.

# scsiblade.lif events

## scsiblade.lif.init.fail

### Severity

ERROR

### Description

This message occurs when the SAN LIF fails to be configured during Vserver initialization. The affected SAN LIF cannot serve data until the issue is resolved.

### Corrective Action

Delete and re-create the network interface by using the "network interface delete -vserver [vserver] -lif [lif-name]" command, and then the "network interface create -vserver [vserver] -lif [lif-name]" command. You can find the configuration information required to create the network interface (data-protocol, home-node, port, etc.) by using the "network interface show -vserver [vserver] -lif [lif-name]" command. If the problem persists, contact NetApp technical support for further assistance.

## Syslog Message

Vserver %s network interface %s type %s could not be initialized (error %u:%u).

## Parameters

**vserver** (STRING): Name of the Vserver.

**lif** (STRING): Name of the network interface (LIF).

**liftype** (STRING): Type of network interface (LIF).

**errortype** (INT): Error type.

**error** (INT): Error.

## scsiblade.lif.init.timeout

### Severity

NOTICE

### Description

This message occurs when Vserver initialization finishes, and some volumes are partially initialized. This can happen when nodes are out of quorum. When the initialization timer expires, LIFs in the Vserver are brought online regardless of whether volume initialization is complete.

### Corrective Action

(None).

## Syslog Message

Delayed LIF initialization timer expired for Vserver "%s".

## Parameters

**vserver** (STRING): Name of the Vserver.

# scsiblade.lu events

## scsiblade.lu.gir.end

### Severity

INFORMATIONAL

### Description

This message occurs when this node completes a group internal rebuild for an LU.

### Corrective Action

(None).

### Syslog Message

Group internal rebuild complete for LUN %s on this node.

### Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.lu.gir.hung

### Severity

NOTICE

### Description

This message occurs when group internal rebuild for an LUN hangs on this node

### Corrective Action

(None).

### Syslog Message

Group internal rebuild hung for LUN %s on this node.

### Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.lu.gir.start

### Severity

INFORMATIONAL

### Description

This message occurs when the distributed LUN group commences an internal rebuild.

### Corrective Action

(None).

## Syslog Message

Group internal rebuild has started for LUN %s on this node.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.lu.incons.memmb

### Severity

ALERT

### Description

This message occurs when the SCSI target identifies a node that has an incorrect view of the group membership for a LUN. The LUN will be unavailable on this node and will have sporadic availability on other nodes.

### Corrective Action

To recover from this condition in HA pair configurations, take over the node with the problem by using the "storage failover takeover" command, and then give back the node by using the "storage failover giveback" command. To recover from this condition in non-HA configurations, reboot the nodes having this condition by using the "system node reboot" command.

## Syslog Message

LUN %s on node %s has an incorrect group membership list. Perform a takeover followed by a giveback of this faulty node.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**nodeName** (STRING): Name of the node that has an incorrect view of the group.

## scsiblade.lu.int.rst.hung

### Severity

ALERT

### Description

This message occurs when the SCSI target cannot complete an internal reset of a LUN. This LUN is unavailable until the nodes that could not complete the internal reset are taken over and given back.

### Corrective Action

To recover from this condition in HA pair configurations, take over the node with the problem by using the "storage failover takeover" command, and then give back the node by using the "storage failover giveback" command. To recover from this condition in non-HA configurations, reboot the nodes having this condition by using the "system node reboot" command.

## Syslog Message

Access to LUN %s is restricted because an internal reset of the LUN was not completed in %d milliseconds. Perform a takeover followed by a giveback for the following nodes: %s.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**lu\_int\_rst\_hung\_time** (INT): Duration of the internal reset attempt in milliseconds.

**nodeName** (STRING): Names of the nodes on which the internal reset of the LUN was not completed.

## **scsiblade.lu.int.rst.pending**

### **Severity**

NOTICE

### **Description**

This message occurs when the SCSI target cannot complete an internal reset of a LUN in 90 seconds. Internal reset is still pending, waiting for internal reset completion within 10 minutes.

### **Corrective Action**

(None).

### **Syslog Message**

Access to LUN %s is restricted because an internal reset of the LUN was not completed in %d milliseconds. Internal reset of the LUN is still pending.

### **Parameters**

**LUNSerialNumber** (STRING): Serial number of the LUN.

**lu\_int\_rst\_pending\_time** (INT): Duration of the internal reset attempt in milliseconds.

**nodeName** (STRING): Names of the nodes on which the internal reset of the LUN was not completed.

## **scsiblade.lu.resync.end**

### **Severity**

NOTICE

### **Description**

This message occurs when the SCSI target internal resynchronization of a LUN is complete.

### **Corrective Action**

(None).

### **Syslog Message**

Resynchronization of LUN %s on this node was completed successfully.

### **Parameters**

**LUNSerialNumber** (STRING): Serial number of the LUN.

## **scsiblade.lu.resync.start**

### **Severity**

NOTICE

### **Description**

This message occurs when the SCSI target internal resynchronization of a LUN starts. The LUN is not accessible on LIFs on this node during the resynchronization, but the LUN will recover automatically.

### **Corrective Action**

(None).

## Syslog Message

Resynchronization of LUN %s has started. The LUN is not accessible on LIFs on this node until resynchronization is complete and the LUN recovers automatically.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

## scsiblade.lu.resync.timeout

### Severity

ALERT

### Description

This message occurs when the SCSI target internal resynchronization of a LUN times out. This LUN is not accessible on this node.

### Corrective Action

To recover from this condition in HA pair configurations, take over the node with the problem by using the "storage failover takeover" command, and then give back the node by using the "storage failover giveback" command. To recover from this condition in non-HA configurations, reboot the nodes having this condition by using the "system node reboot" command.

## Syslog Message

Resynchronization of LUN %s on node %s was not completed in %d seconds. LUN is not accessible on LIFs on this node. Perform a takeover followed by a giveback of this node.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**nodeName** (STRING): Name of the node where resynchronization of the LUN was not completed.

**resync\_time\_out\_time** (INT): Duration of the resynchronization attempt in seconds.

## scsiblade.lu.rtpg.lifs

### Severity

ERROR

### Description

This message occurs when the SCSI Target processes the RTPG command and the number of LIFs per Vserver exceeds 700.

### Corrective Action

Reduce the number of LIFS per Vserver to under 700.

## Syslog Message

The Vserver hosting LUN %s has %d LIFs. RTPG command expects the number of LIFs per Vserver to be below 700.

## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

**num\_lifs** (INT): Number of LIFs on this Vserver.

# scsiblade.lun events

## scsiblade.lun.ids.change

### Severity

ERROR

### Description

This message occurs when the system detects a change in identifiers (serial number, UUID, device ID) of a LUN. As a result, SAN hosts might see the previous identifiers of the LUN. Other LUNs on specified volume are also affected. This can occur due to incomplete processing of a 'snapmirror break' operation for the volume containing this LUN. It can also occur due to the command 'lun rescan' having been issued without following up with bringing the volume offline and then bringing it online.

### Corrective Action

If the LUN is in a mirrored volume and wafiron has been run on the aggregate containing the volume, issue the command 'lun rescan -oovc-to-vol'. Otherwise, bring the volume offline and then online.

### Syslog Message

A change in identifiers of a LUN in Vserver %s and volume with MSID %llu has been detected. The previous and new serial numbers are '%s', '%s'. The previous and new UUIDs are '%s', '%s'. The previous and new device text IDs are '%s', '%s'. The previous and new device binary IDs are '%s', '%s'. Other LUNs on this volume are also affected.

### Parameters

**vserver** (STRING): Name or UUID of the Vserver.  
**volumeMSID** (LONGINT): MSID of the container volume.  
**prevSerial** (STRING): Previous LUN serial number.  
**newSerial** (STRING): New LUN serial number.  
**prevUUID** (STRING): Previous LUN UUID.  
**newUUID** (STRING): New LUN UUID.  
**prevDevTextID** (STRING): Previous LUN Device Text ID.  
**newDevTextID** (STRING): New LUN Device Text ID.  
**prevDevBinID** (STRING): Previous LUN Device Binary ID.  
**newDevBinID** (STRING): New LUN Device Binary ID.

## scsiblade.lun.move.error

### Severity

ERROR

### Description

This message occurs when the result status message of a LUN move job is lost. The LUN move job will not complete until this message is received, even if the LUN has been moved.

### Corrective Action

Use the 'lun move show' command to see if the affected job is still running. If it is, then find the affected LUN path by finding the other EMS entries with the same job UUID and file ID. Once the affected LUN has been found, then offlining and onlining the destination volume to which the LUN is being moved will make the On Demand job manager retry the operation. If the success status is 1 then all lun maps and QoS entries were successfully updated, otherwise they were not. For further assistance, contact NetApp technical support.

## Syslog Message

This message occurs when the result status %d of a LUN move job %s for file ID %llu is lost.

## Parameters

**Success** (INT): Success status contained in the lost message.

**OnDemandUUID** (STRING): UUID of the On Demand job.

**OnDemandIndex** (LONGINT): Index of File in the On Demand Job.

## scsiblade.lun.move.fail

### Severity

ERROR

### Description

This message occurs when a commit operation fails while updating LUN maps during a LUN move operation. The LUN move fails if this occurs.

### Corrective Action

Use the provided reason to help determine the corrective action, and then retry the command. For further assistance, contact NetApp technical support.

## Syslog Message

This message occurs when the RDB update fails for lun move job %s-%llu. Failure reason: %s.

## Parameters

**OnDemandUUID** (STRING): UUID of the On Demand job.

**OnDemandIndex** (LONGINT): Index of File in the On Demand job.

**Reason** (STRING): Reason for the failure.

## scsiblade.lun.offline.system

### Severity

ERROR

### Description

This message occurs when a LUN is brought offline by Data ONTAP® because of a write failure caused by lack of space in the volume.

### Corrective Action

Add space to the volume and bring the LUN online using the "LUN modify" command.

## Syslog Message

LUN %s in volume with MSID %llu on Vserver %s has been brought offline due to lack of space in the volume.

## Parameters

**lunSerial** (STRING): Serial number of the LUN.

**volumeMSID** (LONGINT): MSID of the volume.

**vserver** (STRING): Name or UUID of the Vserver.

## scsiblade.lun.stale.map

### Severity

ERROR

### Description

This message occurs when a LUN or NVMe namespace attributes lookup operation fails because the LUN Vdisk ID or NVMe namespace NSID is not known to Vdisk Operations and Management. This results in degraded switchover, switchback, and Vserver initialization times.

### Corrective Action

Verify that the volume is online by using the "volume show -vserver [vserver name] -volume [volume name] -fields state" command. If necessary, bring the volume online by using the "volume online -vserver [vserver name] -volume [volume name]" command. If you cannot find the volume or the problem persists, contact NetApp technical support for assistance with the diagnostic privilege commands to view and remove stale maps.

### Syslog Message

Could not look up LUN or NVMe namespace attributes for the object with ID '%s' in volume with MSID '%llu' in Vserver '%s'.

### Parameters

**vdiskId** (STRING): Vdisk ID of the LUN or NSID of the NVMe namespace.

**volume** (LONGINHEX): Volume identifier.

**vserver** (STRING): UUID of the Vserver that hosts the LUN or NVMe namespace.

## scsiblade.lun.state.invalid

### Severity

ERROR

### Description

This message occurs when a LUN map operation fails because the state of the LUN in the SAN cache is incorrect.

### Corrective Action

Shut down the node and retry mapping the LUN by issuing the ZAPI/CLI command to one of the other nodes in the cluster. Reboot the node once the LUN has been successfully mapped.

### Syslog Message

The LUN with Vdisk ID %s on Vserver "%s" cannot be mapped because its state in the SAN cache is incorrect.

### Parameters

**vdiskId** (STRING): Vdisk ID of the LUN.

**vserver** (STRING): UUID of the Vserver that hosts the LUN.

## scsiblade.mapped events



## scsiblade.mapped.ns.attr.mod

### Severity

ERROR

### Description

This message occurs when a mapped NVMe namespace attribute changes as a result of a management operation. Changes in NVMe namespace attributes like size, state, or OS type, are not supported in this release.

### Corrective Action

Remove the namespace map using "vserver nvme subsystem map remove" command, then reapply it with the "vserver nvme subsystem map add" command.

### Syslog Message

The Namespace attributes of the Namespace with Vdisk ID '%s' in volume with MSID '%llu' in Vserver '%s' changed.

### Parameters

**vdiskId** (STRING): Vdisk ID of the LUN.  
**volume** (LONGINTEX): Volume identifier.  
**vserver** (STRING): UUID of the Vserver that hosts the LUN.

## scsiblade.mcc events

### scsiblade.mcc.vs.load.wm

### Severity

NOTICE

### Description

This message occurs when the Vserver load during MCC switchover/switchback consumes majority of the 120 second MAX timeout. The LOW and HIGH watermark indicates the possibility of exceeding the 120 second MAX timeout, which could lead to All Paths Down condition on the host.

### Corrective Action

(None).

### Syslog Message

MCC Vserver load hit %s watermark set to %d seconds.

### Parameters

**level** (STRING): watermark level.  
**time** (INT): time.

## scsiblade.mgmt events

### scsiblade.mgmt.evt.lost.mem

**Severity**

ERROR

**Description**

This message occurs when a SAN management event is lost because the system could not allocate enough free memory. Access to LUNs within the Vserver might be affected. The corresponding Vserver might not be available to process any other management changes and events until the node is rebooted to clear this error condition.

**Corrective Action**

Perform a planned takeover and giveback.

**Syslog Message**

A SAN management event was lost because of lack of available memory. Vserver %s might not be available to process additional SAN management commands and events. Access to LUNs within the Vserver might be affected.

**Parameters**

**vserver** (STRING): Name or UUID of the Vserver.

**scsiblade.mgmt.locked****Severity**

NOTICE

**Description**

This message occurs when there is a connectivity loss within the cluster. The SAN configuration is locked and commands to change the SAN configuration are rejected.

**Corrective Action**

(None).

**Syslog Message**

One or more cluster nodes have lost connectivity. The SAN configuration will remain locked and reject any commands to change the configuration until connectivity is restored.

**Parameters**

(None).

**scsiblade.mgmt.unlocked****Severity**

NOTICE

**Description**

This message occurs when connectivity to all cluster nodes is established. The SAN configuration is unlocked and commands to change the SAN configuration are allowed.

**Corrective Action**

(None).

## Syslog Message

Connectivity to all cluster nodes is available. The SAN configuration is unlocked and can be modified.

## Parameters

(None).

## scsiblade.mgmt.wedged

## Severity

EMERGENCY

## Description

This message occurs when the SAN management plane encounters an error that prevents further SAN management commands from being processed.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

SAN management has encountered an error and SAN configuration changes on this node and other nodes in the cluster will fail with timeout errors until corrected.

## Parameters

(None).

## scsiblade.object events

### scsiblade.object.out.of.sync

## Severity

ERROR

## Description

This message occurs when a SAN management event receives a timeout error due to network congestion in the system and all efforts of recovery fail. The specific SAN object for which the management event is intended will be out of sync and will have different values in cache and in persistent store. Access to LUNs within the Vserver might be affected.

## Corrective Action

Contact NetApp technical support. For further information about correcting the problem, see the knowledgebase article 2017369.

## Syslog Message

A SAN management event received a timeout error due to network congestion. Management event %llu on Vserver %s left the SAN object's cache copy out of sync with information in the persistent store. Access to LUNs within the Vserver might be affected.

## Parameters

**opcode** (LONGINT): Operation code.

**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.offline events

## scsiblade.offline

### Severity

ERROR

### Description

This message occurs when the node loses connectivity with the majority of the cluster. When this occurs, both FCP and iSCSI traffic are disabled on this node. Remote cluster traffic is impacted. After connectivity is restored, the node automatically recovers and enables FCP and iSCSI services.

### Corrective Action

Correct the issue that caused the node to lose connectivity with the majority of the cluster. After the issue is corrected, service is restored.

### Syslog Message

This node lost connectivity with the majority of the cluster. When this occurs, both FCP and iSCSI traffic are disabled on this node.

### Parameters

(None).

# scsiblade.online events

## scsiblade.online

### Severity

NOTICE

### Description

This message occurs when the node completes initialization of the SAN configuration cache. FCP and iSCSI traffic are enabled on this node.

### Corrective Action

(None).

### Syslog Message

This node has completed initialization of the SAN configuration cache. FCP and iSCSI traffic are enabled on this node.

### Parameters

(None).

# scsiblade.pport events

## scsiblade.pport.bsh.cancel.timeout

**Severity**

NOTICE

**Description**

PPorts LIF migration started before BSH cancel completed.

**Corrective Action**

No corrective action

**Syslog Message**

PPorts LIF migration started before BSH cancel API completed.

**Parameters**

(None).

**scsiblade.pport.fcp.dbo.err****Severity**

ERROR

**Description**

This message occurs during a nondisruptive operation such as takeover or giveback, when the persistent port feature is not functioning due to configuration errors.

**Corrective Action**

Collect the system logs, and then contact NetApp technical support. Support needs to see the system logs to determine the appropriate corrective action.

**Syslog Message**

FCP LIF "%s" persistence in Vserver "%s" could not be enabled. Error %u/%u: %s.

**Parameters**

**lif** (STRING): Name of the network interface (LIF).

**vserver** (STRING): Name of the Vserver.

**errortype** (INT): Error type.

**error** (INT): Error code.

**errorDecString** (STRING): Error description string.

**scsiblade.pport.fcp.init.err****Severity**

ERROR

**Description**

This message occurs when the persistent port object for a SAN FCP network interface (LIF) fails to be configured during a LIF create operation. LIF persistence will not function during nondisruptive operations like takeover and giveback until the issue is resolved.

**Corrective Action**

Collect the system logs, and then contact NetApp technical support. Support needs to see the system logs to determine the appropriate corrective action.

## Syslog Message

FCP LIF "%s" persistence in Vserver "%s" failed to initialize. Error %u/%u: %s.

## Parameters

**lif** (STRING): Name of the FCP network interface (LIF).

**vserver** (STRING): Name of the Vserver.

**errortype** (INT): Error type.

**error** (INT): Error code.

**errorDecString** (STRING): Error description string.

# scsiblade.prop events

## scsiblade.prop.done.error

### Severity

EMERGENCY

### Description

This message occurs when a CTRAN "proposal done" message is not received. The SAN management plane is inoperative for the Vserver until the corrective action is taken.

### Corrective Action

For further information about correcting the problem, search the knowledgebase of the NetApp technical support web site for the "scsiblade.prop.done.error" keyword.

## Syslog Message

An internal error occurred in a %s group (MSID: %llu). The Vserver %s will not be able to process any management commands.

## Parameters

**groupType** (STRING): CTRAN group type.

**volumeMSID** (LONGINT): MSID of the volume in case the CTRAN group type is "volume"; otherwise, 0.

**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.qosvioldetectregfail events

## scsiblade.QoSViolDetectRegFail

### Severity

ERROR

### Description

This message occurs when the Quality-of-Service (QoS) subsystem fails to register for updates during initialization due to an internal error (e.g. there is not enough system memory available). As a result, workloads might not meet their minimum throughput or their service-level objectives (SLOs).

### Corrective Action

The registration will continue to be retried. If this message continues to occur, perform a planned takeover/giveback sequence to reset the reporting node nondisruptively by using the "takeover -ofnode (reporting\_node)" and "giveback -ofnode (reporting\_node)" commands. If the problem persists, contact

Contact NetApp technical support.

### **Syslog Message**

Failed to register for violation detection updates. QoS min-throughput might operate in a limited capacity.

### **Parameters**

**errorCode** (INT): Error code that was returned during registration.

## **scsiblade.reloading events**

### **scsiblade.reloading**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the local node is recovering from a connectivity loss by purging and reloading the SAN configuration cache. When this occurs, both FCP and iSCSI traffic are temporarily disabled until the reload is complete.

#### **Corrective Action**

(None).

#### **Syslog Message**

This node is reloading the SAN configuration cache to recover from a loss of connectivity. While reloading, both FCP and iSCSI services are disabled. When the reload is complete, both services are automatically re-enabled.

#### **Parameters**

(None).

## **scsiblade.san events**

### **scsiblade.san.config.active**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the SAN active-active state has changed.

#### **Corrective Action**

(None).

#### **Syslog Message**

The symmetric active-active state is %s on %lld LUNs.

#### **Parameters**

**state** (STRING): Active-active state.

**num\_luns** (INT): Number of LUNs.

## **scsiblade.san.rtpgid.err**

### **Severity**

ERROR

### **Description**

This message occurs when a report target port group ID space inconsistency between cluster peers is detected. This can result in no I/O progress or I/O disruption or both on the host.

### **Corrective Action**

Recovery from inconsistency is required. Please contact Contact NetApp technical support. for assistance.

### **Syslog Message**

Consistency group that has LUN %s is in an inconsistent state.

### **Parameters**

**LUNSerialNumber** (STRING): Serial number of the LUN.

## **scsiblade.san.ssm.lu.degraded**

### **Severity**

ALERT

### **Description**

This message occurs when a LUN is marked as degraded due to a SnapMirror® Automated Failover relationship going out of sync. The LUN is no longer accessible from the host.

### **Corrective Action**

Restore LUN access to the host by using the "snapmirror resync" command to bring the out-of-sync relationship back in sync. For assistance, contact NetApp Support.

### **Syslog Message**

LUN %s was marked as degraded due to a SnapMirror Automated Failover relationship going out of sync.

### **Parameters**

**LUNSerialNumber** (STRING): Serial number of the LUN.

## **scsiblade.san.ssm.lu.online**

### **Severity**

NOTICE

### **Description**

This message occurs when a LUN is back online and might be accessible again from the host.

### **Corrective Action**

If the host still can't access the LUN even though the LUN is back online, perform a host-side rescan operation or a reboot to enable the host to access the LUN.

### **Syslog Message**

LUN %s is back online.



## Parameters

**LUNSerialNumber** (STRING): Serial number of the LUN.

# scsiblade.tiered events

## scsiblade.tiered.lun.offline

### Severity

ERROR

### Description

This message occurs when a LUN is brought offline by Data ONTAP® because of too many consecutive failures to access object storage or no response from object storage.

### Corrective Action

Bring object storage online and bring the LUN online using the "LUN modify" command.

### Syslog Message

LUN %s in volume with MSID %llu on Vserver %s has been brought offline due to too many consecutive failed attempts to access object storage or no response from object storage.

## Parameters

**lunSerial** (STRING): Serial number of the LUN.

**volumeMSID** (LONGINT): MSID of the volume.

**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.unavailable events

## scsiblade.unavailable

### Severity

EMERGENCY

### Description

This message occurs when the node loses connectivity with the majority of the cluster and is unable to recover automatically. When this occurs, both FCP and iSCSI traffic are disabled on this node and require manual intervention to come back into service.

### Corrective Action

Reboot the node, and if the node is waiting for giveback, perform a giveback operation.

### Syslog Message

This node lost connectivity with the majority of the cluster. When this occurs, both FCP and iSCSI traffic are disabled and will remain disabled on this node until corrective actions are taken.

## Parameters

(None).

# scsiblade.vol events

## scsiblade.vol.init.failed

### Severity

EMERGENCY

### Description

This message occurs when SAN initialization fails for the specified volume on this node. Access to LUNs contained in this volume from this node is not available.

### Corrective Action

Reboot the node by performing a planned takeover and giveback.

### Syslog Message

SAN configuration in Vserver %s failed to initialize for volume with MSID %llu. Access to LUNs contained in this volume from this node is not available.

### Parameters

**vserver** (STRING): Name or UUID of the Vserver.

**volumeMSID** (LONGINT): MSID of the volume.

**errorType** (LONGINT): Error type.

**errorCode** (LONGINT): Error code.

## scsiblade.vol.lun.stale.map

### Severity

ERROR

### Description

This message occurs when a volume lookup operation fails because the volume is not found. This is an indication that there might be stale LUN maps associated with the volume.

### Corrective Action

Contact NetApp technical support for assistance with the (privilege: diag) "lun stale-map" commands.

### Syslog Message

Could not lookup the volume with MSID '%llu' in Vserver '%s'. This is an indication that there might be stale lun maps associated with the volume.

### Parameters

**volume** (LONGINTEX): Volume identifier.

**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.volume events

## scsiblade.volume.event.lost

### Severity

EMERGENCY

## Description

This message occurs when LUNs in the specified Vserver and volume are not accessible on one or more SAN LIFs. Determine the volume name from an MSID by using the "volume show -vserver [vserver name] -msid [id]" command.

## Corrective Action

The corrective action for this error must be determined from system logs. For more information or assistance, contact NetApp technical support.

## Syslog Message

Access to LUNs in volume with MSID %llu on Vserver %s has been lost through one or more paths.

## Parameters

**volume** (LONGINTEX): Volume identifier (MSID).  
**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.vs events

## scsiblade.vs.purge.fail

### Severity

EMERGENCY

### Description

This message occurs when the local node fails to purge the SAN configuration cache for a Vserver. Manual intervention is required to allow reinitialization of the specified Vserver.

### Corrective Action

Reboot the node, and if the node is waiting for giveback, perform a giveback operation.

### Syslog Message

The local node failed to purge the SAN configuration cache for Vserver %s.

### Parameters

**vserver** (STRING): Name or UUID of the Vserver.

# scsiblade.vserver events

## scsiblade.vserver.op.timeout

### Severity

EMERGENCY

### Description

This message occurs when a command or other internal operation becomes blocked within Data ONTAP®. When this situation occurs, the corresponding Vserver cannot handle any other management changes for SAN until this condition is resolved. Also, the command that timed out might not have been completed, limiting access to LUNs within the Vserver.

### **Corrective Action**

For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "scsiblade.vserver.op.timeout" keyword.

### **Syslog Message**

A Vserver internal operation (0x%08x) has timed out. Vserver %s might not be available to process additional SAN management commands and events.

### **Parameters**

**operation** (LONGINT): Operation code.

**vserver** (STRING): Name or UUID of the Vserver.

# scsitarget events

## scsitarget.fcp events

### scsitarget.fcp.sfp.inserted

#### Severity

INFORMATIONAL

#### Description

This message occurs when a small form-factor pluggable (SFP) is inserted into the Fibre Channel (FC) target adapter.

#### Corrective Action

You can obtain detailed SFP information by using the 'fcp adapter show' command.

#### Syslog Message

An SFP was inserted into FC target adapter %s.

#### Parameters

**adapter** (STRING): Target adapter name.

### scsitarget.fcp.sfp.invalid

#### Severity

ERROR

#### Description

This message occurs when an invalid Small Form-factor Pluggable (SFP) is inserted into the adapter. For example, a 10-GB SFP was inserted into an 8-GB FC adapter.

#### Corrective Action

Replace the SFP with a supported SFP.

#### Syslog Message

The SFP in FC adapter %s is not valid. (%s)

#### Parameters

**adapter** (STRING): Target adapter name.

**reason** (STRING): Reason for invalid SFP.

### scsitarget.fcp.sfp.removed

#### Severity

INFORMATIONAL

#### Description

This message occurs when a Small Form-factor Pluggable (SFP) is removed from the Fibre Channel (FC) target adapter.

## Corrective Action

(None).

## Syslog Message

An SFP was removed from FC target adapter %s.

## Parameters

**adapter** (STRING): Target adapter name.

## scsitarget.fcp.sfp.warning

### Severity

ERROR

### Description

This message occurs when a small form-factor pluggable (SFP) that is not supported for FCoE links is inserted into an adapter that is configured for FCoE. For example, a 1-GB SFP is inserted into a converged network adapter (CNA) port. Note that the underlying Ethernet link might still be supported with the original SFP.

### Corrective Action

Replace the SFP with an SFP that supports FCoE link speed.

### Syslog Message

The SFP in FC adapter %s does not support 10 Gbit/sec speed. FCoE is supported only at 10 Gbit/sec and so the FCoE link will not be online.

### Parameters

**adapter** (STRING): Target adapter name.

## scsitarget.fct events

### scsitarget.fct.linkError

### Severity

ERROR

### Description

This message occurs when a Fibre Channel (FC) Host Bus Adapter (HBA) encounters a physical layer issue. This might result in the frame being dropped. The following are examples of protocol errors that result in this message: - Bad start of frame (SOF) - Bad end of frame (EOF) - CRC error - Running disparity error - Link Reset primitive detection

### Corrective Action

Recurring instances of this message indicate a physical layer issue. Perform the following recovery actions: - Verify that the overall SAN fabric is healthy. - Verify that the SFP is healthy at both endpoints, replace if necessary. - Clean or replace the cables. Contact technical support with the code and status information for further analysis.

## Syslog Message

FC target HBA: %s received unexpected error: %s, %s Status1: 0x%x, Status2: 0x%x, Status3: 0x%x.

## Parameters

**hba** (STRING): Name of the target HBA that issued the error.  
**error1** (STRING): String associated with the error.  
**error2** (STRING): Additional string associated with the error.  
**status1** (INTHEX): Firmware status code 1.  
**status2** (INTHEX): Firmware status code 2.  
**status3** (INTHEX): Firmware status code 3.

## scsitarget.fct.loginfailed

### Severity

ERROR

### Description

This message occurs when the SLI Fibre Channel target driver encounters a FDISC or FLOGI failure. The failure can cause the FC target port or the LIF to be operationally down.

### Corrective Action

Contact technical support for further analysis.

## Syslog Message

FC port %s, NPIV port %d failed fabric login with status of %d and an extended status of x%x.

## Parameters

**hba** (STRING): FC HBA where the login failure occurred.  
**npiv\_port** (INT): NPIV (N Port ID Virtualization) LIF ID.  
**status** (INT): Status value communicated by the firmware after login processing.  
**extended\_status** (INTHEX): Extended status value communicated by the firmware after login processing.

## scsitarget.fct.npivErr

### Severity

ERROR

### Description

This message occurs when the Fibre Channel (FC) fabric does not have NPIV (N Port ID Virtualization) enabled.

### Corrective Action

Enable NPIV on the FC fabric for the FC LIF (logical interface) to come online. The fabric port related to the FC LIF should be disabled and then reenabled after NPIV is enabled. Use the "network interface" command to mark the operational status of the LIF as "up".

## Syslog Message

NPIV is disabled on the fabric %s port %s. The target HBA resides at host %s at port %s.

## Parameters

**switch\_name** (STRING): FC fabric name.

**switch\_port\_num** (STRING): Fabric port number where NPIV is disabled.

**target\_host\_name** (STRING): Host name of the target node.

**target\_host\_port\_name** (STRING): Port name of the target HBA.

## **scsitarget.fct.offlineComplete**

### **Severity**

NOTICE

### **Description**

This message occurs when the FCT adapter offline process is complete.

### **Corrective Action**

(None).

### **Syslog Message**

Fibre Channel target adapter %s was brought offline.

### **Parameters**

**hba** (STRING): HBA (host bus adapter) that was taken offline.

## **scsitarget.fct.offlineStart**

### **Severity**

NOTICE

### **Description**

This message occurs when the FCT adapter offline process starts.

### **Corrective Action**

(None).

### **Syslog Message**

Fibre Channel target adapter %s is being brought offline.

### **Parameters**

**hba** (STRING): Name of the HBA(host bus adapter) that is being taken offline.

## **scsitarget.fct.onlining**

### **Severity**

NOTICE

### **Description**

This message occurs when the FCT adapter is online.

### **Corrective Action**

(None).

### **Syslog Message**

Fibre Channel target adapter %s is being brought online.



## Parameters

**hba** (STRING): Name of the HBA (host bus adapter) that is being brought online.

## scsitarget.fct.pingErr

### Severity

ERROR

### Description

This message occurs when the ELS\_ECHO reply does not match the ELS\_ECHO request issued by the 'fcp ping' CLI command. This means that the 'fcp ping' command failed the signal quality tests specified in the Methodologies for Signal Quality Specification (MSQS).

### Corrective Action

Check the applicable section of MSQS (<http://www.t11.org/ftp/t11/pub/fc/msqs/10-314v3.pdf>): if a default payload was specified in the 'fcp ping' command and the platform is 8GFC or lower then the JSPAT (scrambled jitter pattern) test failed. Check the MSQS section 2.1 for a list of tests that JSPAT verifies. If the platform is 10GFC or higher and a default payload was specified, then an OMA (Optical Modulation Amplitude) or VMA (Voltage Modulation Amplitude) test failed. Check the MSQS section 2.2.1 for an explanation of these methods.

### Syslog Message

FC ping failed because of a mismatch between request and response payloads. The request payload is %s and the response payload is %s at index %d. The link rate of the port is %d.

## Parameters

**req\_payload** (STRING): First mismatched character of ELS ECHO request payload.

**res\_payload** (STRING): First mismatched character of ELS ECHO response payload.

**index** (INT): Index of the first mismatch between the ECHO request and response payloads.

**link\_rate** (INT): Link rate of the Fibre Channel (FC) port.

## scsitarget.fct.port.full

### Severity

NOTICE

### Description

This message occurs when the number of outstanding commands on a physical FC target port exceeds the supported limit on that port.

### Corrective Action

To recover from this condition, evaluate the host fan-in on the port, and then take one of the following actions: - Reduce the number of hosts that are logging in to this port. - Reduce the number of LUs being accessed by the hosts logging in to this port. - Reduce the host command queue depth. Monitor the queue\_full counter on the fcp\_port CM object, and make sure that it is not increasing. Example: statistics show -object fcp\_port -counter queue\_full -instance port.portname -raw Monitor and verify that the threshold counter is also not increasing. Example: statistics show -object fcp\_port -counter threshold\_full -instance port.portname -raw

### Syslog Message

FC target port %s has %d outstanding commands, which exceeds the maximum number of commands %d that can be supported by this port.

## Parameters

**portname** (STRING): Port name on which the outstanding commands exceeded the supported limit.

**active\_commands** (INT): Number of outstanding commands on the port.

**max\_commands** (INT): Number of maximum commands supported on the port.

## scsitarget.fct.port.offline

### Severity

ERROR

### Description

This message occurs when the Fibre Channel target port is taken offline due to too many errors received by the FCT driver. As a result, the port is not able to support FCP traffic.

### Corrective Action

To determine the cause of the errors, run the "fcp adapter show" command and use the output to determine the next course of action.

### Syslog Message

Fibre Channel Target port %s was taken operationally offline due to too many errors received by the FCT driver.

## Parameters

**adapterName** (STRING): The name of the associated Fibre Channel host bus adapter.

## scsitarget.fct.port.thresh

### Severity

NOTICE

### Description

This message occurs when the number of outstanding commands on a physical FC target port exceeds best practices as defined in the SAN configuration guide.

### Corrective Action

To recover from this condition, evaluate the host fan-in on the port, and then take one of the following actions: - Reduce the number of hosts that are logging in to this port. - Reduce the number of LUs being accessed by the hosts logging in to this port. - Reduce the host command queue depth. Monitor the threshold\_full counter on the fcp\_port CM object, and make sure that it is not increasing. Example: statistics show -object fcp\_port -counter threshold\_full -instance port.portname -raw

### Syslog Message

FC target port %s has %d outstanding commands, which exceeds the threshold %d for this port.

## Parameters

**portname** (STRING): Port name on which the outstanding commands exceeded the threshold.

**active\_commands** (INT): Number of outstanding commands on the port.

**supported\_threshold** (INT): Threshold of the number of commands supported on the port.

## scsitarget.fct.portLogin

### Severity

NOTICE

### Description

This message occurs when an initiator logs into a target port.

### Corrective Action

(None).

### Syslog Message

Login at target FC port: '%s' by initiator port: '%s' address 0x%x. The target virtual port is: '%s'.

### Parameters

**hba** (STRING): Name of target Fiber Channel host bus adapter (HBA) where the port logged in.

**portname** (STRING): Port name of the initiator that logged in.

**addr** (INTHEX): Initiator's Fibre Channel 'Port Identifier' or 'Port Address' as defined in FC-PH. This identifier is mapped to the WWPN and reported in the output of the filer CLI "fcp initiator show" command.

**lifname** (STRING): The name of the virtual port of the FC host bus adapter (HBA) where the initiator logged in.

## scsitarget.fct.postFailed

### Severity

EMERGENCY

### Description

This message occurs when the host bus adapter (HBA) failed its power-on self test (POST).

### Corrective Action

Configure the HBA down then up again by using the "fcp config" command. If the HBA consistently fails its POST, replace it. Otherwise, monitor the HBA for intermittent errors and replace it if the errors occur too frequently. As part of the failing POST, an SRAM dump file is saved in the /etc/log directory. Contact NetApp technical support with this file for further analysis.

### Syslog Message

FC target HBA %s could not initialize because of a POST failure.

### Parameters

**hba** (STRING): HBA that failed its POST.

## scsitarget.fct.reset

### Severity

NOTICE

### Description

This message occurs when the FCT adapter is being reset.

## Corrective Action

(None).

## Syslog Message

Resetting Fibre Channel target adapter %s.

## Parameters

**hba** (STRING): Name of the target HBA (host bus adapter) being reset.

## scsitarget.fct.sfpRxPowerLow

### Severity

ALERT

### Description

This message occurs when the power received (RX) by a small form-factor pluggable transceiver (SFP) is at a level below the defined threshold, which might indicate a failing or faulty part.

### Corrective Action

Monitor the operating value. If it continues to decrease, then replace the SFP and/or the cables.

### Syslog Message

The SFP in FC target adapter %s reports that it is receiving (RX) at a low level of power. Operating value %d (uWatts), Threshold value %d (uWatts).

### Parameters

**adapter** (STRING): Target adapter name.

**operating\_value** (INT): The current operating value. The value is in microwatts (uWatts).

**threshold\_value** (INT): The defined threshold value. The value is in microwatts (uWatts).

## scsitarget.fct.sfpTxPowerLow

### Severity

ALERT

### Description

This message occurs when a small form-factor pluggable transceiver (SFP) is transmitting (TX) power at a level below the defined threshold, which might indicate a failing or faulty part.

### Corrective Action

Monitor the operating value. If it continues to decrease then replace the SFP and/or the cables.

### Syslog Message

The SFP in FC target adapter %s reports that it is transmitting (TX) at a low level of power. Operating value %d (uWatts), Threshold value %d (uWatts).

### Parameters

**adapter** (STRING): Target adapter name.

**operating\_value** (INT): The current operating value. The value is in microwatts (uWatts).

**threshold\_value** (INT): The defined threshold value. The value is in microwatts (uWatts).

# scsitarget.hwpfct events

## scsitarget.hwpfct.dump.saved

### Severity

NOTICE

### Description

This message occurs when an unexpected event or illegal condition is detected by the FC target driver and the contents of the FC target's memory is saved. Dump files are stored on the system's root volume in the /etc/log directory, with the following file name format: \*fct\*[*adapter*][date]\_[time].[dmp |queue]

### Corrective Action

Send the dump file to technical support for analysis.

### Syslog Message

A dump for adapter %s was stored in %s.

### Parameters

**adapter** (STRING): Target adapter name.

**filename** (STRING): Name of the file containing the dump.

## scsitarget.hwpfct.els.rjt

### Severity

ERROR

### Description

This message occurs when an Extended Link Services (ELS) request is rejected by the fabric. Depending on the ELS request that is rejected, the initiator might experience issues discovering the target devices.

### Corrective Action

Ensure that the fabric is zoned and configured correctly. Consult NetApp technical support providing the reason code, reason string, and explanation for further analysis.

### Syslog Message

The Fibre Channel ELS request %s sent from adapter %s NPIV port %d to port ID 0x%x failed with reason code 0x%x (%s) and explanation code 0x%x.

### Parameters

**elsrequest** (STRING): ELS request that was rejected by the destination port ID.

**adapterName** (STRING): Name of the associated Fibre Channel host bus adapter (HBA).

**npiv\_port** (INT): NPIV (N Port ID Virtualization) LIF ID of the source of the ELS request.

**portID** (INHEX): Destination port ID of the ELS request.

**reasonCode** (INHEX): Reason code provided by the destination port ID for rejection of the ELS request.

**reason** (STRING): Reason text provided by the destination port ID for rejection of the ELS request.

**explanationCode** (INHEX): Explanation code provided by the destination port ID qualifying the stated reason.

## scsitarget.hwpfct.errorReset

### Severity

NOTICE

### Description

This message occurs when the FCT driver encounters an error which was resolved by resetting the adapter.

### Corrective Action

(None).

### Syslog Message

An error was encountered in the FC target driver on Fibre Channel target adapter %s. The adapter will be automatically reset to clear the %s condition.

### Parameters

**hba** (STRING): Name of the target HBA (host bus adapter) to reset.

**details** (STRING): Details of the error that caused the reset.

## scsitarget.hwpfct.extraHBA

### Severity

ERROR

### Description

This message occurs when the system detects more than the maximum number of supported Fibre Channel target host bus adapters (HBAs). The extra HBAs are not initialized.

### Corrective Action

Remove the extra Fibre Channel target HBAs from the system.

### Syslog Message

HBA %s will not be initialized. More than %d Fibre Channel target HBAs were found in the system.

### Parameters

**hba** (STRING): Name of the Fibre Channel target HBA (host bus adapter) that were not initialized.

**max** (INT): Maximum number of supported Fibre Channel target HBAs.

## scsitarget.hwpfct.fabRegFail

### Severity

ALERT

### Description

This message occurs when Fibre Channel target port fails to register with the fabric. As a result the intended initiators can not see the target devices.

### Corrective Action

If fabric did not send a rejection the next step is to check cables and connection with the switch. Also ensure that the fabric is zoned correctly. Consult NetApp technical support providing the details and if

available, the reason code, string and explanation for further analysis.

### **Syslog Message**

Fibre Channel target adapter %s unable to register fc4\_type with fabric and export target devices.

### **Parameters**

**hba** (STRING): The name of the target HBA (host bus adapter) that failed.

**portID** (INT): Port ID of the target port trying to register with the fabric.

**command** (INT): Command that failed during fabric registration.

**reasonCode** (INTHEX): Reason code (as defined in FC-FS) of the rejection response from the fabric.

**reasonString** (STRING): Reason string explaining why the request from target was rejected.

**explanation** (INTHEX): Explanation code (as defined in FC-FS). This provides further detail on the rejection.

## **scsitarget.hwpfct.globalLogout**

### **Severity**

NOTICE

### **Description**

This message occurs when the FC HBA (host bus adapter) detects a link down condition and all the initiators have logged out.

### **Corrective Action**

Confirm that the filer is still connected to the switch and that the switch is functional.

### **Syslog Message**

FC Target %s: Link down, all initiators logged out.

### **Parameters**

**hba** (STRING): FC target adapter name.

## **scsitarget.hwpfct.initFailed**

### **Severity**

ALERT

### **Description**

This message occurs when the FCT driver fails to initialize the HBA (host bus adapter).

### **Corrective Action**

Replace the adapter.

### **Syslog Message**

Initialization of adapter failed on Fibre Channel target adapter %s.

### **Parameters**

**hba** (STRING): (None).

## scsitarget.hwpfct.linkUp

### Severity

NOTICE

### Description

This message occurs when FCT driver receives a link up notification.

### Corrective Action

(None).

### Syslog Message

Link up on Fibre Channel target adapter %s.

### Parameters

**hba** (STRING): Name of the target HBA which received link-up.

## scsitarget.hwpfct.thirdPartyLogout

### Severity

NOTICE

### Description

This message occurs when the FCT driver determines that an initiator has issued a third-party logout. Because the initiator is logged out, all SCSI commands from this port are aborted. If the global bit was set, all initiators are logged out and all commands are terminated.

### Corrective Action

Check to ensure that the initiator is still connected to the fabric.

### Syslog Message

FC target %s: The initiator at port ID: 0x%x (WWPN %s) issued a third-party logout with %s.

### Parameters

**hba** (STRING): FC target adapter name.

**port\_id** (INT): Initiator's Fibre Channel "Port Identifier" or "Port Address", as defined in FC-PH. This identifier is mapped to the WWPN and reported in the output of the system CLI command "fcp initiator show" command.

**node** (STRING): WWPN of the initiator that changed and had commands reset.

**globalBit** (STRING): Message to indicate whether the global bit in the third-party logout frame was set.

## scsitarget.hwpfct.tooManyInitiatorsConnected

### Severity

ERROR

### Description

This message occurs when the system detects too many initiators on the Fibre Channel fabric.

### Corrective Action

Remove initiators from the fabric until the maximum is no longer being hit. You can remove initiators via zoning or physical disconnection.



## Syslog Message

Fibre Channel target adapter %s cannot have more than %d initiators connected. Resetting adapter.

## Parameters

**hba** (STRING): System name for the target HBA(host bus adapter) that has too many initiators attempting to connect to it.

**max\_initiators** (INT): Max initiators that can connect to the HBA.

# scsitarget.ispfct events

## scsitarget.ispfct.AENErr

### Severity

ALERT

### Description

This message occurs when Fibre Channel host bus adapter(HBA) returns an event notification to the driver but the driver did not recognize the event type.

### Corrective Action

Contact technical support with the code and status information for further analysis.

## Syslog Message

Fibre Channel target HBA %s received unexpected AEN code 0x%x, Status1 0x%x, Status2 0x%x, Status3 0x%x and Status4 0x%x.

## Parameters

**hba** (STRING): Name of the target HBA that gave the error.

**code** (INHEX): Firmware error code.

**status1** (INHEX): Firmware status code 1.

**status2** (INHEX): Firmware status code 2.

**status3** (INHEX): Firmware status code 3.

**status4** (INHEX): Firmware status code 4.

## scsitarget.ispfct.cfmodeChanged

### Severity

NOTICE

### Description

This event generates a notification that the cfmode was changed. Changing the cfmode should only be done in coordination with FCP connected hosts as this can cause devices to change, appear, or disappear on the hosts.

### Corrective Action

(None).

## Syslog Message

FCP cfmode changed from %s to %s.

## Parameters

**old\_mode** (STRING): The cfmodesetting before the change.

**new\_mode** (STRING): The cfmodesetting after the change.

## scsitarget.ispfct.configChanged

### Severity

NOTICE

### Description

This message occurs when the ISPFCT driver determines that a port configuration has changed.

### Corrective Action

Check to ensure the initiator is still connected to the fabric.

### Syslog Message

FCP target %s reported a port configuration change. %s.

## Parameters

**hba** (STRING): FCP target adapter name.

**message** (STRING): Description of the configuration change.

## scsitarget.ispfct.fwErr

### Severity

ERROR

### Description

This message occurs when firmware on the Fibre Channel (FC) target host bus adapter (HBA) returns a system error. The system is resetting the HBA in an attempt to recover from the error, but further analysis is necessary to determine the cause of the problem and whether further corrective action is needed.

### Corrective Action

Contact NetApp technical support with the subsequent SRAM dump so that the cause of the problem and applicable corrective action can be determined.

### Syslog Message

Firmware on FC target HBA %s returned a system error code 0x%x, status1 0x%x, status2 0x%x, status3 0x%x, and status4 0x%x. Resetting HBA to recover from the error, but SRAM dump analysis is needed.

## Parameters

**hba** (STRING): Name of the target HBA that failed.

**code** (INHEX): Firmware error code.

**status1** (INHEX): Firmware status code 1.

**status2** (INHEX): Firmware status code 2.

**status3** (INHEX): Firmware status code 3.

**status4** (INHEX): Firmware status code 4.

## scsitarget.ispfct.fwfileChecksumFailed

**Severity**

ALERT

**Description**

This message occurs when the Fibre Channel driver fails to load the firmware file when initializing the adapter. The adapter cannot serve data until the firmware is loaded successfully.

**Corrective Action**

Reboot the node to reinitialize the adapter. If the adapter still fails to initialize, replace the adapter. For further assistance, contact NetApp technical support.

**Syslog Message**

Fibre Channel target adapter %s firmware file %s failed checksum after reading %d bytes.

**Parameters**

**hba** (STRING): Adapter that failed to load the firmware file.

**filename** (STRING): Name of file that could not be loaded.

**bytesread** (INT): Number of bytes that have been read in from the firmware file so far.

**scsitarget.ispfct.invalidNVRAMNodeName****Severity**

ALERT

**Description**

This message occurs when ONTAP® software reads an invalid World Wide Node Name (WWNN) from a Fibre Channel target adapter's NVRAM. A default WWNN will be temporarily assigned to the adapter. This might cause fabric zoning not to work correctly and might cause problems with anything that depends on consistent WWNNs and WWPNS.

**Corrective Action**

Reboot the node. If the problem still exists, replace the adapter. Contact NetApp technical support for assistance.

**Syslog Message**

Invalid WWNN of Fibre Channel target adapter %s. Assigning temporary default: %s.

**Parameters**

**hba** (STRING): Name of the target HBA that has invalid NVRAM

**default\_name** (STRING): Node name being applied to the HBA.

**scsitarget.ispfct.linkBreak****Severity**

ERROR

**Description**

ispfct link break notification

**Corrective Action**

If link break is unexpected check cables, switches and the other host (if direct connected). Please check the

Firmware status code 1, 2 and 4 for reason.

### **Syslog Message**

Link break detected on Fibre Channel target adapter %s. Firmware status code status1 0x%x, status2 0x%x, and status4 0x%x.

### **Parameters**

**hba** (STRING): Name of the target HBA which has lost link  
**status1** (INHEX): Firmware status code 1.  
**status2** (INHEX): Firmware status code 2.  
**status4** (INHEX): Firmware status code 4.

## **scsitarget.slifct events**

### **scsitarget.slifct.asyncError**

#### **Deprecated**

Deprecated as of ONTAP® version 9.2.0 in an effort reduce event messages.

#### **Severity**

ERROR

#### **Description**

This message occurs when the Fibre Channel target device driver encounters an unrecognized asynchronous event notification from the firmware, possibly as a result of a firmware update. The unrecognized event is ignored by the system.

#### **Corrective Action**

No administrative action is required. If the messages persist, contact NetApp technical support with the event details for assistance.

### **Syslog Message**

FC port %s encountered an unexpected event of x%x, type of x%x, with additional status of x%x, x%x, x%x.

### **Parameters**

**hba** (STRING): Name of the port where the event occurred.  
**async\_event\_code** (INHEX): Asynchronous event code that was unrecognized by the target driver.  
**async\_event\_type** (INHEX): Asynchronous event type that was unrecognized by the target driver.  
**async\_event\_word1** (INHEX): Asynchronous event word\_1 that was unrecognized by the target driver.  
**async\_event\_word2** (INHEX): Asynchronous event word\_2 that was unrecognized by the target driver.  
**async\_event\_word3** (INHEX): Asynchronous event word\_3 that was unrecognized by the target driver.

## **scsitarget.slifct.fwAutoBoot**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the SLIFCT driver successfully updates the firmware version of the host bus adapter (HBA). The controller will be rebooted automatically to complete the process.

### Corrective Action

None; the controller will be rebooted automatically to enable the new HBA firmware.

### Syslog Message

FC target HBA %s firmware version updated from %s to %s. The controller will be rebooted automatically to enable the new HBA firmware.

### Parameters

**hba** (STRING): HBA where the firmware version update occurred.  
**currentversion** (STRING): Firmware version the HBA was running.  
**bundleversion** (STRING): Firmware version bundled with Data ONTAP®.

## scsitarget.slifct.fwNeedBoot

### Severity

ERROR

### Description

This message occurs when the SLIFCT driver successfully updates the firmware version of the host bus adapter (HBA), and requires a reboot of the controller to complete the process.

### Corrective Action

Reboot the controller to enable the new HBA firmware.

### Syslog Message

FC target HBA %s firmware version updated from %s to %s. Reboot the controller to enable the new HBA firmware.

### Parameters

**hba** (STRING): HBA where the firmware version update occurred.  
**currentversion** (STRING): Firmware version the HBA was running.  
**bundleversion** (STRING): Firmware version bundled with Data ONTAP®.

## scsitarget.slifct.fwUpdDone

### Severity

NOTICE

### Description

This message occurs when the SLIFCT driver successfully updates the firmware version of the host bus adapter (HBA).

### Corrective Action

(None).

### Syslog Message

FC target HBA %s firmware version updated from %s to %s.

### Parameters

**hba** (STRING): HBA where the firmware version update occurred.  
**currentversion** (STRING): Firmware version the HBA was running.

**bundleversion** (STRING): Firmware version bundled with Data ONTAP®.

## **scsitarget.slifct.fwUpdErr**

### **Severity**

ERROR

### **Description**

This message occurs when the SLIFCT driver encounters an error while updating the firmware version of the host bus adapter (HBA).

### **Corrective Action**

Contact technical support for further analysis.

### **Syslog Message**

FC target HBA %s failed to update the firmware version from %s to %s with the status of %d.

### **Parameters**

**hba** (STRING): HBA where the firmware version update failure occurred.

**currentversion** (STRING): Firmware version the HBA is running.

**bundleversion** (STRING): Firmware version bundled with Data ONTAP®.

**status** (INT): Status code of the firmware update process.

## **scsitarget.slifct.linkBreak**

### **Severity**

ERROR

### **Description**

This message occurs when the SLIFCT driver receives a link break notification.

### **Corrective Action**

If the link break was unexpected, check the cables and switch connection.

### **Syslog Message**

Link break detected on Fibre Channel target HBA %s with event status %d , topology type %d, status1 0x%x, status2 0x%x.

### **Parameters**

**hba** (STRING): Name of the target HBA (Host bus adapter) that lost link.

**status** (INHEX): Reported event's status.

**topology** (INHEX): Code for identifying the topology.

**status1** (INHEX): Value of the SLIPOINT\_ERROR1 register.

**status2** (INHEX): Value of the SLIPOINT\_ERROR2 register.

## **scsitarget.slifct.mboxError**

### **Deprecated**

Deprecated as of ONTAP® version 9.2.0 in an effort to reduce event messages.

**Severity**

ERROR

**Description**

This message occurs when the SLIFCT host bus adapter (HBA) encounters an error while processing a mailbox command.

**Corrective Action**

Contact NetApp technical support with the command, subsystem, subopcode, status, and extended status code for further analysis.

**Syslog Message**

FC port %s has a mailbox processing failure for command x%x, subsystem x%x, subopcode x%x, with the status of %d and the extended status of x%x.

**Parameters**

**hba** (STRING): HBA where the mailbox command failure occurred.

**command** (INTEX): Mailbox command code.

**subsystem** (INTEX): Mailbox subsystem code.

**subopcode** (INTEX): Mailbox command subsystem opcode.

**status** (INT): Status communicated by the firmware for the mailbox command.

**extended\_status** (INTEX): Extended status communicated by the firmware for the mailbox command.

**scsitarget.slifct.rebootRequired****Severity**

EMERGENCY

**Description**

This message occurs when a Fibre Channel adapter experiences an unrecoverable error. The adapter cannot serve data without a node reboot.

**Corrective Action**

Reboot the node to reinitialize the adapter. If the adapter fails to initialize or continues to experience errors, replace the adapter. For further assistance, contact NetApp technical support.

**Syslog Message**

Unrecoverable error detected on Fibre Channel target HBA %s with error codes %d and %d.

**Parameters**

**hba** (STRING): Name of the target HBA that encountered the error.

**ERROR\_1** (INT): Value of the SLIPOINT\_ERROR1 register.

**ERROR\_2** (INT): Value of the SLIPOINT\_ERROR2 register.

**scsitarget.slifct.retryExc****Deprecated**

Deprecated as of ONTAP® version 9.2.0 in an effort to reduce event messages.

**Severity**

ERROR

## Description

This message occurs when the SLI Fibre Channel Target driver encounters multiple Work Queue Entry (WQE) failures of a certain type that exhaust the retry limit.

## Corrective Action

Contact NetApp technical support with the type, command code, subtype, status, and extended status values for further analysis.

## Syslog Message

FC port %s, NPIV port %d has reached the retry limit %d of WQE processing failure for type %s, command %s, subtype x%x with status of %d and an extended status of x%x.

## Parameters

**hba** (STRING): FC HBA where the WQE failure occurred.

**npiv\_port** (INT): NPIV LIF ID.

**retry\_limit** (INT): Number of times that this WQE can be retried.

**type** (STRING): Type of WQE.

**command** (STRING): Name of the command in the WQE.

**subtype** (INTHEX): Subtype of the command.

**status** (INT): Status value communicated by the firmware after WQE processing.

**extended\_status** (INTHEX): Extended status value communicated by the firmware after WQE processing.



# secd events

## secd.asidtrans events

### secd.asidTrans.invalidAsid

#### Severity

ALERT

#### Description

This message occurs when ONTAP® software cannot translate an abbreviated SID to a Windows SID. An abbreviated SID is an internal ID maintained by ONTAP. This failure can cause CIFS share access issues.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

For Vserver "%s", ONTAP could not translate an abbreviated SID to a Windows SID.

#### Parameters

**vserverName** (STRING): Vserver associated with this operation.

## secd.authsys events

### secd.authsys.lookup.failed

#### Severity

ERROR

#### Description

This message occurs when the incoming UNIX user ID (UID) that tries to mount or access a mount point cannot be looked up in any of the name-services (NIS, LDAP, file).

#### Corrective Action

Ensure that the UNIX UID that is trying to mount or access a mount point is part of a name-service, such as NIS, LDAP, or file.

#### Syslog Message

Unable to retrieve credentials for UNIX user with UID (%u) on Vserver (%s) for client with IP address (%s).

#### Parameters

**uid** (INT): UNIX UID.

**vserverName** (STRING): Name of the Vserver.

**clientIP** (STRING): IP Address of the client.

## secd.cifs events

## secd.cifs.machacct.missing

### Severity

ALERT

### Description

This message occurs when the CIFS server machine account is missing from the Active Directory domain.

### Corrective Action

Create the missing CIFS server machine account in Active Directory under the appropriate domain and organizational unit. Then run "vserver cifs password-change" command to begin using it.

### Syslog Message

The CIFS server machine account '%s' for Vserver '%s' is missing from the Active Directory domain '%s'.

### Parameters

**acctName** (STRING): Name of the missing CIFS server machine account.

**vserverName** (STRING): Vserver whose CIFS server machine account is missing.

**domainName** (STRING): Name of the Active Directory domain.

## secd.cifsauth events

### secd.cifsAuth.denied

#### Severity

ERROR

#### Description

This message occurs when a CIFS authentication attempt fails for any reason other than an unknown user name or bad password.

#### Corrective Action

Examine the failure details to determine corrective action. A common failure is the inability to communicate with a domain controller.

#### Syslog Message

vserver (%s) Cannot authenticate CIFS user. %s

#### Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

### secd.cifsAuth.noNameMap

#### Severity

ERROR

#### Description

This message occurs when a CIFS authentication succeeds, but system cannot map the Windows user to an appropriate UNIX user.

## Corrective Action

Examine the failure details to determine corrective action. Common failures include no appropriate Windows-to-UNIX name mapping rules, no configured default UNIX user, or the inability of the system to communicate with NS-SWITCH authorization sources.

## Syslog Message

vserver (%s) CIFS name to UNIX name mapping problem. %s

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.cifsAuth.noUnixCreds

### Severity

ERROR

### Description

This message occurs when the system is unable to retrieve necessary UNIX credentials for a Windows user during the CIFS authentication process.

### Corrective Action

Examine the failure details to determine corrective action. A common failure is the inability to communicate with NS-SWITCH authorization sources.

### Syslog Message

vserver (%s) Cannot get UNIX credentials for CIFS user. %s

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.cifsAuth.problem

### Severity

ERROR

### Description

This message occurs when a CIFS authentication attempt fails for any reason other than an unknown user name or bad password.

### Corrective Action

Examine the failure details to determine corrective action. Common failures include the inability to communicate with domain controllers, NIS servers, or LDAP servers due to connectivity or configuration problems.

### Syslog Message

vserver (%s) General CIFS authentication problem. %s

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

# secd.conn events

## secd.conn.auth.failure

### Severity

NOTICE

### Description

This message occurs when the Vserver cannot establish a TCP/UDP connection to or be authenticated by an outside server such as NIS, LSA, LDAP and KDC. Subsequently, some features of the storage system relying on this connection might not function correctly.

### Corrective Action

Ensure that the server being accessed is up and responding to requests. Ensure that there are no networking issues stopping the Vserver from communicating with this server. If the error reported is related to an authentication attempt, ensure that any related configurable user credentials are set correctly.

### Syslog Message

Vserver (%s) could not %s over the network to server (%s)%s. Error: %s (%s).

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

**connAuthAttempt** (STRING): Connection or authentication attempt.

**serverInfo** (STRING): Name or IP:Port of the server that was related to the failure.

**sourceInfo** (STRING): Information about the source that was making the connection or authentication.

**ErrnoMessage** (STRING): Errno message from the connection or authentication failure.

**operation** (STRING): Operation being performed.

# secd.dc events

## secd.dc.out.of.pipe.instances

### Severity

ALERT

### Description

This message occurs when the Windows Domain Controller (DC), while responding to an MSRPC request from Data ONTAP®, returns an error indicating that it ran out of pipe instances in the listening state. This can happen when the requests from Data ONTAP are received by the DC during a timing window when the limited number of listening pipe instance are already allocated to other requests. Such errors received by Data ONTAP might result in longer CIFS/NFS authentication and/or access times, potentially resulting in CIFS/NFS client timeouts.

### Corrective Action

Contact Microsoft support.

## Syslog Message

MSRPC for Vserver %s to %s:%s was retried %d times (delay: %d usecs).

## Parameters

**vserver** (STRING): Vserver that encountered the failure.

**dc** (STRING): DC that returned the error.

**pipe** (STRING): Pipe on the DC.

**retries** (INT): Number of retries.

**delay** (INT): Time (usecs) spent retrying.

# secd.dns events

## secd.dns.server.timed.out

### Severity

ERROR

### Description

This message occurs when the DNS server fails to respond to a query and timeout occurs.

### Corrective Action

Make sure that the DNS server is up and running and that there are no networking issues preventing the Vserver from communicating with the DNS server.

## Syslog Message

DNS server %s did not respond to vserver = %s within timeout interval.

## Parameters

**ipaddr** (STRING): Ip address of the DNS server.

**vsName** (STRING): Name of the Vserver for which DNS response timed out.

## secd.dns.srv.lookup.failed

### Severity

ERROR

### Description

This message occurs when the Domain Name Service (DNS) server fails to look up a service name.

### Corrective Action

Make sure that the service name is input correctly and that there are no networking issues preventing the Vserver from communicating with the DNS server.

## Syslog Message

DNS server failed to look up service (%s) for vserver (%s) with error (%s).

## Parameters

**serviceName** (STRING): Service name to look up.

**vsName** (STRING): Name of the Vserver.

**errorMsg** (STRING): Error message of the failure.

# secd.ipq events

## secd.ipq.invalid.entry

### Severity

ERROR

### Description

This message occurs when the IP qualified name-mapping entry is omitted because of an error during the name-mapping check. For example, an error can occur if a host name is configured as the value for ip-qualifier but the hostname lookup fails, in combination with no IP address defined in the name-mapping configuration to validate against the incoming user.

### Corrective Action

After resolving the error reported in this message, run the "vserver name-mapping refresh-hostname-ip" command to refresh the name-mapping information.

### Syslog Message

Skipping the name-mapping entry on Vserver "%s" for user "%s" at position %u. Reason: %s.

### Parameters

**vserverName** (STRING): Vserver associated with this operation.

**userName** (STRING): Name of the user.

**position** (INT): Name-mapping position.

**errorString** (STRING): Description of the specific error condition: for example, "Failed to lookup hostname gpo.co.in."

# secd.kerberos events

## secd.kerberos.clockskew

### Severity

ERROR

### Description

This message occurs when there is a "time error"(clock skew, time skew, time out of bounds). This error indicates that there is a time discrepancy between client and node or client and Key Distribution Center (KDC). The Kerberos authentication request from the client or the node was forwarded to the KDC and it failed because the timestamp encrypted in the Kerberos ticket was different by more than the maximum time difference that is configured on the KDC.

### Corrective Action

Ensure that the clock time of the node is identical to that of the client and to that of the KDC. To keep the node and KDC time clocks in synchronization automatically, configure Network Time Protocol (NTP) services on the node. Increasing the clock skew interval may also alleviate this condition: To do so, modify the Kerberos-realm configuration clock-skew parameter (denoted as "Maximum tolerance for computer clock synchronization" in Windows® Active Directory) from the default 300 seconds to 600 seconds or more. Note: Increasing the clock-skew interval makes the client protocols less secure against network replay attacks.

## Syslog Message

Kerberos client or node clock skew error for vserver (%s)%s

## Parameters

**vserverName** (STRING): Name of the vserver that is having the error.

**clientInfo** (STRING): Information of the client that is having the error if available.

## secd.kerberos.lookupFailed

### Severity

ERROR

### Description

This message occurs when the Kerberos user is not a part of any name-service. ONTAP maps the Kerberos Service Principal Name(SPN) to a NFS user name while establishing a security context. If the NFS user name is not found in any of the name-services (LDAP, NIS, file), it leads to a failure in establishing security context, which in turn fails the Kerberos mount.

### Corrective Action

Ensure that there is a corresponding UNIX user name for the Kerberos Service Principal Name(SPN) in name services such as NIS, LDAP, or file.

## Syslog Message

Unable to map Kerberos user (%s) to appropriate UNIX user on Vserver (%s).

## Parameters

**uname** (STRING): Kerberos NFS user.

**vserverName** (STRING): Name of the Vserver.

## secd.kerberos.noAuthdata

### Severity

ERROR

### Description

This message occurs when a Kerberos ticket for a user does not contain authorization data.

### Corrective Action

Ask the user to obtain a new, valid Kerberos ticket and map the share again.

## Syslog Message

Kerberos client has no authorization data for Vserver "%s" with user account "%s".

## Parameters

**vserverName** (STRING): Name of the Vserver on which the error occurred.

**userAccount** (STRING): User account for which there is insufficient credential information.

## secd.kerberos.preauth

**Severity**

ERROR

**Description**

This message occurs when invalid credentials are provided for an Active Directory user or the machine account password is out of sync with the credentials set in the Active Directory.

**Corrective Action**

If the reported error is due to invalid credentials, make sure valid credentials are provided for the Active Directory user. Otherwise, if the reported error is due to an out-of-sync machine account password, run the "vserver cifs password-reset -vserver vserver\_name" command to update the password in the Active Directory.

**Syslog Message**

A Kerberos pre-authentication failure occurred for SVM "%s" due to %s.

**Parameters**

**vserverName** (STRING): Name of the storage VM (SVM) that is having the error.

**errReason** (STRING): Reason for the pre-authentication error.

**secd.kerberos.tktexpired****Severity**

ERROR

**Description**

This message occurs when the client's ticket has expired. This error indicates that the timestamp encrypted in the client's Kerberos ticket has exceeded its maximum lifetime or expired.

**Corrective Action**

Ensure that the clock time of the node is identical to that of the client and to that of the KDC. To keep the node and KDC time clocks in synchronization automatically, configure Network Time Protocol (NTP) services on the node. Increasing the clock skew interval may also alleviate this condition: To do so, modify the Kerberos-realm configuration clock-skew parameter (denoted as "Maximum tolerance for computer clock synchronization" in Windows® Active Directory) from the default 300 seconds to 600 seconds or more. Note: Increasing the clock-skew interval makes the client protocols less secure against network replay attacks.

**Syslog Message**

Kerberos client ticket has expired for vserver (%s)%s

**Parameters**

**vserverName** (STRING): Name of the vserver that is having the error.

**clientInfo** (STRING): Information of the client that is having the error if available.

**secd.kerberos.tktnyv****Severity**

ERROR



## Description

This message occurs when the client presented a ticket to the server that is not yet valid (in relationship to the server time). This error indicates that the clocks on the KDC and the client are not synchronized.

## Corrective Action

Ensure that the clock time of the node is identical to that of the client and to that of the KDC. To keep the node and KDC time clocks in synchronization automatically, configure Network Time Protocol (NTP) services on the node. Increasing the clock skew interval may also alleviate this condition: To do so, modify the Kerberos-realm configuration clock-skew parameter (denoted as "Maximum tolerance for computer clock synchronization" in Windows® Active Directory) from the default 300 seconds to 600 seconds or more. Note: Increasing the clock-skew interval makes the client protocols less secure against network replay attacks.

## Syslog Message

Kerberos client ticket not yet valid for vserver (%s)%s

## Parameters

**vserverName** (STRING): Name of the vserver that is having the error.

**clientInfo** (STRING): Information of the client that is having the error if available.

# secd.ldap events

## secd.ldap.bindDn.missing

### Severity

ALERT

### Description

This message occurs when the `-bind-dn` user is not configured and the `-bind-as-cifs-server` option is set to false in the Lightweight Directory Access Protocol (LDAP) client configuration. This LDAP configuration is not recommended because it might result in CIFS credentials being used for authentication if CIFS configuration is present (even though `-bind-as-cifs-server` is not set). It might also result in anonymous authentication attempts.

### Corrective Action

Configure the bind distinguished name (user) and the bind user's password using the following commands:  
`vserver services name-service ldap client modify -bind-dn vserver services name-service ldap client modify-bind-password`

### Syslog Message

The LDAP bind DN user is not configured on Vserver %s.

### Parameters

**vserverName** (STRING): Vserver associated with this operation.

## secd.ldap.connectFailure

### Severity

ALERT

## Description

This message occurs when the server could not establish a TCP connection to a Lightweight Directory Access Protocol (LDAP) server.

## Corrective Action

From a LDAP client workstation, make sure that the LDAP server is responding to requests. Also make sure that the portmapper on the LDAP server is responding to requests. Make sure that there are no networking issues stopping the cluster from communicating with this LDAP server.

## Syslog Message

vserver (%s) could not make a connection over the network to LDAP server (%s) at address (%s) and received error (%s)

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**serverName** (STRING): Name of the LDAP server that was not responding.

**serverAddress** (STRING): Address of the LDAP server that was not responding.

**ldapererror** (STRING): Internal LDAP client library error.

## secd.ldap.hostnames.not.resolved

### Severity

ERROR

### Description

This EMS is generated when none of the configured LDAP hostnames can be resolved.

### Corrective Action

Examine the DNS configuration for the corresponding Vserver.

### Syslog Message

None of the hostnames can be resolved for Vserver: %s.

### Parameters

**vserverName** (STRING): Vserver associated with this operation.

## secd.ldap.hostnames.resolved.partially

### Severity

ERROR

### Description

This EMS is generated when some of the configured LDAP hostnames cannot be resolved.

### Corrective Action

Examine the DNS configuration for the corresponding Vserver.

### Syslog Message

Hostnames cannot be resolved for Vserver: %s, unresolvedHosts: %s.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**unresolvedHostnames** (STRING): Hostnames which cannot be resolved.

## secd.lldap.noServers

### Severity

EMERGENCY

### Description

This message occurs when none of the configured Lightweight Directory Access Protocol (LDAP) servers are accepting connections.

### Corrective Action

From an LDAP client workstation, make sure that all configured LDAP servers are responding to requests. Ensure that there are no networking issues stopping the cluster from communicating with the configured LDAP servers. Also, ensure that the portmapper running on the LDAP server is working correctly.

### Syslog Message

None of the LDAP servers configured for Vserver (%s) are currently accessible via the network for LDAP service type (%s).

## Parameters

**vserverName** (STRING): Vserver associated with this operation.  
**ldapOperation** (STRING): LDAP operation and service for which the connection is required.

## secd.lldap.query.timed.out

### Severity

ERROR

### Description

This message occurs when the LDAP server fails to respond to a query and timeout occurs.

### Corrective Action

Because LDAP server timeouts might be the result of connectivity issues between the storage controller and the external service or latency in LDAP server response time, ensure that external connectivity and external services have not been disrupted.

### Syslog Message

Vserver '%s': LDAP server %s did not respond to query within timeout (%d seconds) interval.

## Parameters

**vsName** (STRING): Name of the Vserver for which LDAP response timed out.  
**ipaddr** (STRING): IP address of the LDAP server.  
**timeout** (INT): Number of seconds before the LDAP query is timed out.

## secd.lldap.referralError

**Severity**

INFORMATIONAL

**Description**

This message indicates that the domain controller does not have a copy of the requested object (which exists) and it is providing a location that is more likely to hold the object.

**Corrective Action**

Modify the environment to avoid LDAP referrals.

**Syslog Message**

Server (%s) does not hold the target entry for ldap filter (%s) on vserver (%s).

**Parameters**

**serverAddress** (STRING): Address of the LDAP server that responded with referral.

**ldapSearchFilter** (STRING): The Ldap search filter associated with this operation.

**vserverName** (STRING): Vserver associated with this operation.

**secd.ldap.sasl.bind.delayed****Severity**

ERROR

**Description**

This message occurs when a Lightweight Directory Access Protocol (LDAP) server responds slowly to SASL bind requests. This might result in longer SMB/CIFS authentication times, potentially resulting in SMB/CIFS client timeouts. Also, this might affect other LDAP bind requests and Security Daemon (SecD) delays.

**Corrective Action**

Ensure that there are no networking issues creating intermittent communication problems with the LDAP server. Make sure that the machine running LDAP is responsive and not overloaded.

**Syslog Message**

LDAP SASL bind taking longer time on server "%s" for SVM "%s".

**Parameters**

**serverAddress** (STRING): Address of the LDAP server that is not responding fast enough.

**vserverName** (STRING): SVM associated with this operation.

**secd.ldap.slowServer****Severity**

ERROR

**Description**

This message indicates that the Lightweight Directory Access Protocol (LDAP) server is not responding to requests in the expected time frame.

**Corrective Action**

Make sure that there are no networking issues creating intermittent communication problems with the LDAP

server. Make sure that the machine running LDAP is responsive and not overloaded. From a LDAP client workstation, run LDAP commands to verify long response times.

### Syslog Message

from CIFS Server(%s) calls to LDAP server (%s) at address (%s) is executing slowly enough to adversely impact the performance of your server.

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**serverName** (STRING): Name of the LDAP server that was not responding fast enough.

**serverAddress** (STRING): Address of the LDAP server that was not responding fast enough.

## secd.ldap.starttls.delayed

### Severity

ERROR

### Description

This message occurs when Lightweight Directory Access Protocol (LDAP) server responses to STARTTLS requests take 5 or more seconds. This might result in longer SMB/CIFS authentication times, potentially resulting in SMB/CIFS client timeouts. Also, this might affect other LDAP bind requests and Security Daemon (SecD) delays.

### Corrective Action

Ensure that there are no networking issues creating intermittent communication problems with the LDAP server. Make sure that the machine running LDAP is responsive and not overloaded.

### Syslog Message

The STARTTLS operation is taking too long on LDAP server "%s" for SVM "%s".

### Parameters

**serverAddress** (STRING): Address of the LDAP server that is not responding fast enough.

**vserverName** (STRING): SVM associated with this operation.

## secd.lsa events

### secd.lsa.noServers

### Severity

EMERGENCY

### Description

This message occurs when none of the configured LSA servers are accepting connections.

### Corrective Action

Ensure that all configured LSA servers are responding to requests. Ensure that there are no networking issues stopping the cluster from communicating with the configured LSA servers.

### Syslog Message

None of the LSA servers configured for Vserver (%s) are currently accessible via the network.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

# secd.nameTrans.events

## secd.nameTrans.groupNotFound

### Severity

ERROR

### Description

This message occurs when the group name cannot be resolved by any of the entities in the ns-switch for the virtual server.

### Corrective Action

If this is a Windows® name, make sure that the name exists in Active Directory. If this is a UNIX® name, make sure that it is in one of the configured ns-switch entities for the virtual server. Make sure that it was entered locally as a unix-group, in the Network Information Service (NIS) maps, or in the Lightweight Directory Access Protocol (LDAP) server.

### Syslog Message

vserver (%s) could not resolve the group name (%s) by any of the entities in the ns-switch for the virtual server.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**groupName** (STRING): Name of the group that could not be converted to an ID.

## secd.nameTrans.invalidConfig

### Severity

ERROR

### Description

This message occurs when the name resolution configuration for the virtual server is not set up correctly.

### Corrective Action

Consult the documentation and correct the ns-switch entry for the virtual server.

### Syslog Message

vserver (%s) has an incorrectly configured ns-switch. Check all of the virtual servers to ensure that the ns-switch has been configured correctly.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

## secd.nameTrans.invalidUser

### Severity

ERROR

## Description

This message occurs when the user name cannot be resolved by any of the entities in the ns-switch for the virtual server.

## Corrective Action

If this is a Windows® name, make sure the name exists in Active Directory. If this is a UNIX® name, make sure it is in one of the configured ns-switch entities for the virtual server. Make sure that the name was entered locally as a unix-user, in the Network Information Service (NIS) maps, or in the Lightweight Directory Access Protocol (LDAP) server.

## Syslog Message

vserver (%s) could not resolve user name (%s).

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**userName** (STRING): Name of the user that could not be converted to an ID.

## secd.nameTrans.noNameMapping

### Severity

ERROR

### Description

This message occurs when no name mappings or default-users are defined to convert a Windows® name to a UNIX® name or a UNIX name to a Windows name.

### Corrective Action

Make sure that a default Windows user is defined for the Network File System (NFS) configuration. Make sure that a default UNIX user is defined for the CIFS configuration. Decide to either to create one of the previously mentioned items, or create a mapping rule that maps the user or group in the correct direction.

### Syslog Message

vserver (%s) could not map name (%s). Reason: %s.

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**userName** (STRING): Name of the user that could not be converted to an ID.

**attemptedMapping** (STRING): Description of the specific error condition.

## secd.nameTrans.unknownUser

### Severity

ERROR

### Description

This message occurs when the user name cannot be resolved by any of the entities in the ns-switch for the virtual server.

### Corrective Action

If this is a Windows® name, make sure that the name exists in Active Directory. If this is a UNIX® name, make sure that it is in one of the configured ns-switch entities for the virtual server. Make sure that it was entered locally as a unix-user, in the Network Information Service (NIS) maps, or in the Lightweight

Directory Access Protocol (LDAP) server.

### **Syslog Message**

vserver (%s) could not resolve the user name (%s) by any of the entities in the ns-switch for the virtual server

### **Parameters**

**vserverName** (STRING): vserver associated with this operation.

**userName** (STRING): Name of the user that could not be converted to an ID.

## **secd.nameTrans.userNotFound**

### **Severity**

ERROR

### **Description**

This message occurs when the user name cannot be resolved by any of the entities in the ns-switch for the virtual server.

### **Corrective Action**

If this is a Windows® name, make sure that the name exists in Active Directory. If this is a UNIX® name, make sure it is in one of the configured ns-switch entities for the virtual server. Make sure that it was entered locally as a unix-user, in the Network Information Service (NIS) maps, or in the Lightweight Directory Access Protocol (LDAP) server.

### **Syslog Message**

vserver (%s) could not resolve the user name (%s) by any of the entities in the ns-switch for the virtual server.

### **Parameters**

**vserverName** (STRING): vserver associated with this operation.

**userName** (STRING): Name of the user that could not be converted to an ID.

## **secd.netgroup events**

### **secd.netgroup.ldap.badFilter**

### **Severity**

ERROR

### **Description**

This message occurs if the filter for searching the Lightweight Directory Access Protocol (LDAP) server is found to be invalid. The typical reason for this issue is incorrect LDAP client configuration or a bad netgroup name. Searching for the current netgroup is skipped, so all hosts in the netgroup might not be authorized.

### **Corrective Action**

Use the "vserver services name-service ldap client schema show" command to ensure that the Object Class and Attribute in the LDAP client schema are correct. Check the LDAP configuration on the server side to ensure that the attribute value (such as nisNetgroup) is correct.



## Syslog Message

LDAP search on vserver (%s) failed on invalid filter: %s.

## Parameters

**vserverName** (STRING): The name of the Vserver associated with this operation.

**filter** (STRING): The invalid filter that caused the LDAP search to fail.

# secd.netlogon events

## secd.netlogon.noServers

### Severity

EMERGENCY

### Description

This message occurs when none of the configured Netlogon servers are accepting connections.

### Corrective Action

Ensure that all configured Netlogon servers are responding to requests. Ensure that there are no networking issues stopping the cluster from communicating with the configured Netlogon servers.

## Syslog Message

None of the Netlogon servers configured for Vserver (%s) are currently accessible via the network.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

# secd.nfs events

## secd.nfs.groupLimit

### Severity

ERROR

### Description

This message occurs when Data ONTAP® truncates the number of groups to which a user belongs. This can happen when a user of a UNIX-based system is in more groups than the limit supported by the system.

### Corrective Action

Determine the maximum number of groups allowed for the Vserver in question with the "nfs show -vserver <vservname> -fields extended-groups-limit" command. If necessary, configure the limit of extended groups on the Vserver by using the (privilege mode: advanced) command "nfs modify -vserver <vsname> -extended-groups-limit <0..1024>". Make sure that the user does not belong to more than the maximum number of groups supported on the site NIS or LDAP server. (This maximum should correspond to the value of the "extended-groups-limit" field that was either displayed or set in the given commands.)

## Syslog Message

User (%s) is in too many groups (%d): on Vserver (%s).

## Parameters

**uname** (STRING): UNIX user name.

**gidCount** (INT): Total number of auxillary groups that a user belongs to, in various name services.

**vserverName** (STRING): Name of the Vserver.

# secd.nfsauth events

## secd.nfsAuth.noCifsCred

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails because of the inability of the system to retrieve a matching CIFS credential for use in multi-protocol security operations.

### Corrective Action

Examine the failure details to determine corrective action. This failure usually occurs because the system is unable to communicate with Active Directory.

### Syslog Message

vserver (%s) NFS authorization cannot retrieve CIFS credentials. %s

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.nfsAuth.noCifsSid

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails because the system could not resolve an associated Windows security identifier (SID).

### Corrective Action

Examine the failure details to determine corrective action. This failure usually occurs because the system is unable to communicate with Active Directory.

### Syslog Message

vserver (%s) Unable to resolve CIFS security identifier (SID). %s

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.nfsAuth.noCifsUser

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails because the Windows user that NFS maps to cannot be found.

### Corrective Action

Examine the failure details to determine corrective action. Common failures include faulty UNIX-to-Windows name mapping rules, an improperly configured default Windows user, or the system is unable to communicate with Active Directory.

### Syslog Message

vserver (%s) Mapped CIFS user name not found. %s

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.nfsAuth.noNameMap

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails because of a UNIX to Windows name mapping issue.

### Corrective Action

Examine the failure details to determine corrective action. Common failures include no appropriate UNIX-to-Windows name mapping rules, no configured default Windows user, or the inability of the system to contact LDAP if LDAP is configured for name mapping.

### Syslog Message

vserver (%s) Cannot map UNIX name to CIFS name. %s

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.nfsAuth.noUnixCreds

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails because credentials for a UNIX user cannot be determined.

## Corrective Action

Examine the error journal for the collection of events that led to the specific failure. A common failure is the inability to communicate with NS-SWITCH authorization sources.

## Syslog Message

Vserver "%s" cannot determine UNIX identity. %s

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem.

## secd.nfsAuth.problem

### Severity

ERROR

### Description

This message occurs when an NFS authorization attempt fails.

### Corrective Action

Examine the failure details to determine corrective action. Common failures include name mapping issues, or the inability to communicate with domain controllers, NIS servers, or LDAP servers due to connectivity or configuration problems.

### Syslog Message

vserver (%s) General NFS authorization problem. %s

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**failureDetail** (STRING): Error journal is a collection of events leading up to a failure. This failure is likely to impact many (if not all) users, and might reflect a misconfiguration or environmental problem..

## secd.nis events

### secd.nis.connectFailure

#### Deprecated

Deprecated as of version 9.5 since this EMS is no longer needed.

#### Severity

ALERT

#### Description

This message occurs when the server could not establish a TCP connection to a Network Information Service (NIS) server.

#### Corrective Action

From a UNIX ® workstation, make sure that the NIS server is responding to requests. Also make sure that the portmapper on the NIS server is responding to requests. Make sure that there are no networking issues stopping the cluster from communicating with this NIS server.

## Syslog Message

vserver (%s) could not make a connection over the network to NIS server (%s) at address (%s) and received error (%s)

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**serverName** (STRING): Name of the NIS server that was not responding.

**serverAddress** (STRING): Address of the NIS server that was not responding.

**ErrnoMessage** (STRING): Errno message from the connect failure.

## secd.nis.noServers

### Deprecated

Deprecated as of version 9.5 since this EMS is no longer needed.

### Severity

EMERGENCY

### Description

This message occurs when none of the configured Network Information Service (NIS) servers are accepting connections.

### Corrective Action

From a UNIX® workstation, make sure that all configured NIS servers are responding to requests. Ensure that there are no networking issues stopping the cluster from communicating with the configured NIS servers. Also, ensure that the portmapper running on the NIS server is working correctly.

## Syslog Message

None of the NIS servers configured for Vserver (%s) are currently accessible via the network.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

## secd.nis.slowServer

### Deprecated

Deprecated as of version 9.5 since this EMS is no longer needed.

### Severity

ERROR

### Description

This message indicates that the Network Information Service (NIS) server is not responding to requests in the expected time frame.

### Corrective Action

Make sure that there are no networking issues creating intermittent communication problems with the NIS server. Make sure that the machine running NIS is responsive and not overloaded. From a UNIX ® workstation, run YP commands to verify long response times.

## Syslog Message

From CIFS Server(%s) calls to NIS server (%s) at address (%s) is executing slowly enough to adversely impact the performance of your server.

## Parameters

**vserverName** (STRING): vserver associated with this operation.

**serverName** (STRING): Name of the NIS server that was not responding fast enough.

**serverAddress** (STRING): Address of the NIS server that was not responding fast enough.

# secd.nonetgroupfile events

## secd.noNetgroupFile

### Severity

ALERT

### Description

This message occurs when trying to process a netgroup in an export policy while the Vserver "ns-switch" option is set exclusively to "file" but no netgroup file is loaded. Client requests cannot be processed until a netgroup file is loaded or another name service is added to the Vserver "ns-switch" option.

### Corrective Action

Load a netgroup file for the Vserver by using the "netgroup load" command, or add another name service to the Vserver "ns-switch" option.

## Syslog Message

Cannot evaluate a netgroup in an export policy without a properly configured "ns-switch" option for Vserver "%s". It is set exclusively to "file" but no netgroup file is loaded.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

# secd.quark events

## secd.quark.ddns.updt.failure

### Severity

ERROR

### Description

This message occurs when the server can not a send dynamic DNS (DDNS) update to the DNS server.

### Corrective Action

For "NOTIMP" errors, enable DDNS on the DNS server. For "REFUSED" errors, verify the DDNS settings on the DNS server. For all other errors, Contact NetApp technical support..

## Syslog Message

DDNS update failed with error %s.

## Parameters

**cmdError** (STRING): Error from the attempted DDNS update.

## secd.quark.sid.lookup.failed

### Severity

ERROR

### Description

This message occurs when the Active Directory server cannot find the Security Identifier (SID) for a given account name.

### Corrective Action

Provide a valid account name.

### Syslog Message

Failed to lookup the SID for account name "%s" with error "%s".

## Parameters

**accountName** (STRING): Name of the account whose SID is looked up in the server.

**errorMsg** (STRING): Error received from the server while doing a SID lookup for the given account name.

## secd.rpc.events

### secd.rpc.authRequest.blocked

#### Severity

ALERT

#### Description

This message occurs when a CIFS authentication RPC is blocked by Security Daemon(SecD) due to continuous authentication requests with wrong logon password from a particular client.

#### Corrective Action

Check whether there is an unexpected increase in the number of authentication requests with wrong logon password from a particular CIFS client.

#### Syslog Message

Too many CIFS authentication attempts with wrong password from client "%s" on Vserver "%s".

## Parameters

**clientIP** (STRING): Client that has been blocked.

**vserverName** (STRING): Vserver associated with this operation.

### secd.rpc.server.ready

#### Severity

INFORMATIONAL

## Description

This message occurs when SecD successfully loads the configuration and is ready to serve all RPCs.

## Corrective Action

(None).

## Syslog Message

SecD is ready to serve all RPCs.

## Parameters

(None).

## secd.rpc.server.request.dropped

### Severity

ERROR

### Description

This message occurs when a remote procedure call(RPC) is dropped by SecD due to a lack of memory space. It might occur either when SecD receives too many requests than it can handle or when a number of requests have accumulated due to server connectivity issues.

### Corrective Action

Check whether there is an unexpected increase in the number of authentication requests from NFS/CIFS clients. Verify that there are no connectivity issues to the external servers like DNS, LDAP, and DC.

### Syslog Message

RPC "%s" that was sent from "%s" was dropped by SecD due to a lack of memory space.

### Parameters

**rpcName** (STRING): Name of the RPC that has been dropped.

**callerName** (STRING): Caller of the RPC that has been dropped.

## secd.single events

### secd.single.label.domain

#### Severity

ERROR

#### Description

This message occurs when the trusted domains for the CIFS home domain of a Vserver contains a single-label domain (SLD). An SLD configuration is not supported because it can cause CIFS access issues.

#### Corrective Action

Examine and remove any SLDs configured for the CIFS domain associated with the corresponding Vserver on your domain controller. Additionally, examine and revoke any account's access to the CIFS server where the account is from the trusted SLD.



## Syslog Message

Single-label domain (SLD) "%s" was found in the CIFS trusted domains for Vserver %s.

## Parameters

**domainName** (STRING): Trusted SLD.

**vserverName** (STRING): Vserver associated with this operation.

# secd.strong events

## secd.strong.auth.required

### Severity

ALERT

### Description

This message occurs when the external LDAP server configured for a Vserver is enforcing stronger authentication than what the ONTAP® LDAP client is configured for.

### Corrective Action

Configure stronger LDAP session security using the following commands: 1. vserver services name-service ldap client modify -session-security 2. vserver cifs security modify -session-security-for-ad-ldap

## Syslog Message

Stronger authentication enforced by LDAP server for Vserver %s.

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

# secd.unexpectedfailure events

## secd.unexpectedFailure

### Severity

ERROR

### Description

This message occurs when the security daemon captures an unexpected failure.

### Corrective Action

Examine the failure details to determine corrective action.

## Syslog Message

Unexpected SecD failure in Vserver "%s". Details: %s

## Parameters

**vserverName** (STRING): Vserver associated with this operation.

**failureDetail** (STRING): The error journal documents the events leading up to the failure. This failure is likely to impact many (if not all) users, and it might indicate a misconfiguration or environmental problem.

# secd.unixlookupfailure events

## secd.unixLookupFailure

### Severity

ERROR

### Description

This message occurs when an UNIXi® user or group lookup fails. This failure is likely to impact many or all users, and might indicate a misconfiguration or environmental problem.

### Corrective Action

Examine the error journal in the /mroot/etc/mlog/secd.log file to learn more about why the LOOKUP failed. The error journal is a collection of events leading up to a failure. Common failures include the wrong user name or ID and the inability to communicate with the NIS or LDAP servers due to connectivity or configuration problems. If the vserver's NSSWITCH is set to "file", verify that the user or group that is being looked up has been created using the cluster's ngshell interface. If the vserver's NSSWITCH is set to "NIS", run 'yp' commands from an NIS client workstation to verify that the UNIX lookup works on the configured NIS server. If the vserver's NSSWITCH is set to "LDAP", run 'ldap' commands from an LDAP client workstation to verify that the UNIX lookup works on the configured LDAP server. Verify that the cluster can communicate with the configured NIS or LDAP servers.

### Syslog Message

UNIX lookup failure on Vserver (%s) for client with IP address (%s). %s

### Parameters

**vserverName** (STRING): vserver associated with this operation.

**clientIP** (STRING): IP Address of the client.

**failureDetail** (STRING): Reason for the failure.

# sectrace events

## sectrace.warning events

### sectrace.warning.aboutperf

#### Severity

ALERT

#### Description

This message occurs when a security tracing filter is enabled. It can cause a significant performance degradation and should be disabled as soon as it is no longer needed.

#### Corrective Action

Consider disabling the Security tracing filters as soon as it is no longer needed. 'sectrace filter modify -vserver [id] -index [index] -enabled disabled'

#### Syslog Message

Security tracing is enabled on Vserver %s for the next %d minutes.

#### Parameters

**vservername** (STRING): Name of the vserver on which file/directory tracing is enabled.

**duration** (INT): Duration of the sectrace filter that is enabled.

# securepurge events

## securepurge.aborted events

### securepurge.aborted

**Severity**

ERROR

**Description**

This message occurs when a secure purge operation is aborted in a flexible volume.

**Corrective Action**

Correct the cause of the error, and then retry the operation.

**Syslog Message**

Secure purge operation is aborted on the volume '%s' in Vserver '%s'. Error: '%s'.

**Parameters**

**volumename** (STRING): Name of the volume.

**vservername** (STRING): Name of the Vserver.

**desc** (STRING): Reason for the failure.

## securepurge.completed events

### securepurge.completed

**Severity**

NOTICE

**Description**

This message occurs when a secure purge operation is completed in a flexible volume.

**Corrective Action**

(None).

**Syslog Message**

Secure purge operation is completed on the volume '%s' in Vserver '%s'.

**Parameters**

**volumename** (STRING): Name of the volume.

**vservername** (STRING): Name of the Vserver.

## securepurge.started events

### securepurge.started

**Severity**

NOTICE

**Description**

This message occurs when a secure purge operation is started in a flexible volume.

**Corrective Action**

(None).

**Syslog Message**

Secure purge operation is started on the volume '%s' in Vserver '%s'.

**Parameters**

**volumename** (STRING): Name of the volume.

**vservername** (STRING): Name of the Vserver.

# security events

## security.invalid events

### security.invalid.login

#### Severity

ALERT

#### Description

This message occurs when an attempt is made to access the appliance by using invalid authentication credentials.

#### Corrective Action

Verify that the user made a genuine attempt to access the system. If not, use the "security login delete" command to either remove the account or use the "security login lock" command to lock the account to disallow access to the system using that user account. Also, strengthen the authentication credentials of the user account by updating it. More detail regarding the authentication failure is available in audit.log file.

#### Syslog Message

Failed to authenticate login attempt to Vserver: %s, username: %s, application: %s.

#### Parameters

**vserver** (STRING): The name of the Vserver to which the access was attempted.

**userName** (STRING): The user name used to gain access to the Vserver.

**application** (STRING): The name of the application used to gain access to the Vserver.

# service events

## service.resp events

### service.resp.committed

#### Severity

NOTICE

#### Description

This message occurs when a service (Firefly, Slice Service, etc.) is detected as responsive and this state is committed into the Zookeeper database. The service ID and node ID on which it is running are displayed in the event notification.

#### Corrective Action

(None).

#### Syslog Message

The responsiveness for service (ID: %u) of type %s running on node %s (ID: %u, Uuid: %s) has been committed into the database.

#### Parameters

**serviceID** (INT): Unique identifier (service ID) of the service.

**serviceType** (STRING): Type of the service.

**nodeName** (STRING): Name of the node on which the service is running.

**nodeID** (INT): Unique identifier of the node on which the service is running.

**nodeUuid** (STRING): Universal unique identifier of the node on which the service is running.

## service.unresp events

### service.unresp.committed

#### Severity

NOTICE

#### Description

This message occurs when a service (Firefly, Slice Service, etc.) is detected as unresponsive and this state is committed into the Zookeeper database. The service ID and node ID on which it is running are displayed in the event notification.

#### Corrective Action

(None).

#### Syslog Message

The unresponsiveness for service (ID: %u) of type %s running on node %s (ID: %u, Uuid: %s) has been committed into the database.

#### Parameters

**serviceID** (INT): Unique identifier (service ID) of the service.

**serviceType** (STRING): Type of the service.

**nodeName** (STRING): Name of the node on which the service is running.

**nodeID** (INT): Unique identifier of the node on which the service is running.

**nodeUuid** (STRING): Universal unique identifier of the node on which the service is running.



# ses events

## ses.access events

### ses.access.noEnclServ

#### Severity

ALERT

#### Description

This message occurs when Enclosure Services in the storage system cannot establish contact with the enclosure monitoring process in any disk shelf on the channel. Enclosure Services in the storage system tries periodically to reestablish contact with the disk shelf.

#### Corrective Action

If the error persists for more than an hour, then reboot the controller.

#### Syslog Message

No Enclosure Services detected through channel %s.

#### Parameters

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

### ses.access.noMoreValidPaths

#### Severity

ALERT

#### Description

This message occurs when Enclosure Services in the storage system loses access to the enclosure monitoring process in the disk shelf I/O module.

#### Corrective Action

Enclosure Services in the storage system tries periodically to reestablish contact with the disk shelf. If the error persists for more than an hour, then halt the storage system, power-cycle the disk shelf, and reboot. If the error still persists, then the shelf I/O module might need to be replaced.

#### Syslog Message

No more valid paths to Enclosure Services in shelf %d on channel %s.

#### Parameters

**shelf** (INT): Disk shelf number on the channel with which access has been lost.

**channelName** (STRING): Name of the adapter port or switch port associated with the shelf; for example, "7a" or "myswitch:5".

### ses.access.noShelfSES

**Severity**

ERROR

**Description**

This message occurs when Enclosure Services in the storage system cannot establish contact with the Enclosure Services process in the indicated disk shelf. The system tries periodically to reestablish contact with the disk shelf.

**Corrective Action**

If the error persists for more than an hour, then the disk shelf module must be power-cycled. If the module is a system CPU module (such as IOM6E), or if the system is attached in a single path to the module, then halt the system, reseal the module, and reboot. In a multipath system, remove the shelf module, wait 30 seconds, and then reinsert the module. If the error persists, then replace the shelf module.

**Syslog Message**

No SCSI-3 Enclosure Services on channel %s shelf %d.

**Parameters**

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

**shelf** (INT): Disk shelf number on the channel with which contact has been lost.

**ses.access.sesUnavailable****Severity**

ALERT

**Description**

This message occurs when Enclosure Services in the storage system cannot establish contact with the enclosure monitoring process in one or more disk shelves on the channel. The system tries periodically to reestablish contact with the disk shelves.

**Corrective Action**

If the error persists for more than an hour, then reboot the controller.

**Syslog Message**

Enclosure Services unavailable for one or more shelves on channel %s.

**Parameters**

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

**ses.badsharestorageconfigerr events****ses.badShareStorageConfigErr****Severity**

EMERGENCY

## Description

This message occurs when a disk shelf module that is not supported in a Shared Storage system, such as an LRC module, is detected in a Shared Storage system.

## Corrective Action

Replace the unsupported module with one that is supported.

## Syslog Message

Disk shelf %d on adapter %s has an unsupported module for a Shared Storage configuration.

## Parameters

**shelf\_num** (INT): Disk shelf number the unsupported module is on.

**ha\_name** (STRING): Host adapter on which the disk shelf resides.

# ses.bridge events

## ses.bridge.fw.getFailWarn

### Severity

ERROR

### Description

This message occurs when the bridge firmware revision cannot be obtained.

### Corrective Action

Check the connection to the bank of Maxtor drives.

### Syslog Message

On channel %s, the bridge firmware revision could not be obtained using disk IDs %d-%d.

### Parameters

**ChannelName** (STRING): Name of the adapter or switch port that the system tried to get the revision number of.

**DiskID** (INT): Starting ID of the bank of Maxtor disks that was used to obtain the bridge firmware.

**DiskRange** (INT): Number of drives in that bank of Maxtor disks.

## ses.bridge.fw.mmErr

### Severity

ERROR

### Description

This message occurs when the bridge firmware revision is inconsistent.

### Corrective Action

Check the firmware revision number and make sure that they are consistent. You might have to update the firmware.

### Syslog Message

On channel %s, the bridge firmware revision appears to be inconsistent : %8.8s and %8.8s.

## Parameters

**ChannelName** (STRING): Name of the channel that contains the inconsistent firmware revision.

**PrevRev** (STRING): Previous firmware revision.

**CurRev** (STRING): Current firmware revision.

# ses.channel events

## ses.channel.rescanInitiated

### Severity

INFORMATIONAL

### Description

This message occurs when a rescan for enclosure services devices is initiated.

### Corrective Action

(None).

### Syslog Message

Initiating rescan on channel %s.

## Parameters

**channelName** (STRING): Name of the adapter port or switch port being rescanned; for example, "7a" or "myswitch:5".

# ses.config events

## ses.config.diskStackMaxError

### Severity

ERROR

### Description

This message occurs when a stack of shelves contains more disk bays than the maximum number that are allowed.

### Corrective Action

Remove shelves from the storage stack until the the number of disk bays do not exceed the maximum number.

### Syslog Message

There are a total of %d disk bays on shelves attached to channel "%s", which exceeds the maximum number of disk bays (%d).

## Parameters

**numDiskBays** (INT): Number of disk bays on the storage stack.

**channelName** (STRING): Name of the adapter port or switch port connected to the storage stack associated with the error (for example, "7a" or "myswitch:5").

**maxDiskBays** (INT): Maximum number of disk bays that are allowed on the storage stack.

## ses.config.drivePopError

### Severity

ERROR

### Description

This message occurs when the channel has more disk drives on it than are allowed. Systems using synchronous mirroring allow more disk drives per channel than other systems.

### Corrective Action

If synchronous mirroring is intended to be used, ensure that the license is installed. Otherwise, reduce the number of disk drives on the channel to no more than the maximum allowed.

### Syslog Message

Disk drive population is greater than %d in channel %s; this is an illegal configuration. Synchronous mirroring %s is enabled.

### Parameters

**fcalTenancy** (INT): Maximum number of disk drives allowed on this loop.

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

**isIsNot** (STRING): Whether synchronous mirroring is enabled.

## ses.config.IllegalEsh270

### Severity

EMERGENCY

### Description

This message occurs when ONTAP® software detects one or more Embedded Switched Hub (ESH) disk shelf modules in a disk shelf that is attached to a FAS270. This is not a supported configuration.

### Corrective Action

Replace the ESH modules with ESH2 modules.

### Syslog Message

One or more ESH modules are attached to a FAS270. This is not a supported configuration. Replace the ESH modules with ESH2 modules.

### Parameters

(None).

## ses.config.shelfMixError

### Severity

EMERGENCY

### Description

This message occurs when the channel has a mixture of ATA and Fibre Channel disk shelves; this is not a supported configuration.

### Corrective Action

Mixed-mode operation of ATA and Fibre Channel disks on the system is only supported on separate loops. Move all Fibre Channel-based disk shelves to one loop and place all Fibre Channel-to-ATA-based disk shelves on another loop.

### Syslog Message

Disk shelves present on channel %s are of different drive types; this is not a supported configuration.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

## ses.config.shelfPopError

### Severity

ERROR

### Description

This message occurs when the channel has more disk shelves on it than are allowed.

### Corrective Action

Reduce the number of disk shelves on the channel to the number specified.

### Syslog Message

Disk shelf population on channel %s is %d. Reduce the number of disk shelves on this channel to %d or fewer.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port associated with the error; for example, "7a" or "myswitch:5".

**numShelves** (INT): Number of disk shelves on the port.

**maxShelves** (INT): Maximum number of disk shelves on a port.

## ses.disk events

### ses.disk.configOk

### Severity

INFORMATIONAL

### Description

This message occurs when there are no longer any drives in a FAS2050 system's slots 20 through 23.

### Corrective Action

(None).

### Syslog Message

Disk configuration errors have been cleared.

**Parameters**

(None).

**ses.disk.illegalConfigWarn****Severity**

ERROR

**Description**

This message occurs when disk drives are inserted into the bottom row of a FAS2050 storage appliance. Disk drives are not supported in those slots.

**Corrective Action**

Remove the disk drives.

**Syslog Message**

Disk drives are inserted into unsupported drive slots in the bottom row of the storage appliance.

**Parameters**

(None).

**ses.download events****ses.download.powerCyclingChannel****Severity**

INFORMATIONAL

**Description**

This message occurs when the power cycling channel event is issued after a disk shelf firmware download to disk shelves that require a power-cycle to activate the new code.

**Corrective Action**

(None).

**Syslog Message**

Power cycling channel %s.

**Parameters**

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf is being power-cycled.

**ses.download.shelfToReboot****Severity**

INFORMATIONAL

**Description**

This message occurs after the completion of shelf firmware transfer to the DS14-Mk2-AT (R200) disk shelf, at which point the disk shelf requires about another five minutes to transfer the new firmware to its nonvolatile program memory, whereupon it reboots to begin to execute the new firmware. During this

reboot, a Fibre Channel loop reinitialization occurs, temporarily interrupting the loop.

#### **Corrective Action**

(None).

#### **Syslog Message**

The DS14-Mk2-AT shelves will be rebooted shortly to run the new firmware. The firmware download process should take about three minutes per shelf.

#### **Parameters**

(None).

## **ses.download.suspendIOForPowerCycle**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the suspending I/O event signals that the storage subsystem is temporarily stopping I/O to disks while one or more disk shelves have their power cycled after a download, if required by the disk shelf design.

#### **Corrective Action**

(None).

#### **Syslog Message**

Suspending disk I/O prior to shelf power cycle.

#### **Parameters**

(None).

## **ses.drive events**

### **ses.drive.badPlacement**

#### **Severity**

ALERT

#### **Description**

This message occurs when disk drives are placed in locations that do not provide for proper signal integrity. In some older shelf types only, circuitry in the disk drives conditions the signals on their way through the disk shelves. If the disk drives are not in certain locations, the signal might degrade and cause errors. Disk drives must be kept in the proper locations even if one or more disk drives are removed for maintenance reasons.

#### **Corrective Action**

If only loop A is used, always keep at least one disk drive in bays 0-2 and at least one disk drive in bays 3-6. If loop B is also used, always keep at least one disk drive in bays 0-2 and at least one disk drive in bays 4-6. The integrity requirement of both loops is satisfied if disk drives are placed as required for loop B. Remember that bays 3 and 4 are used for enclosure services access in XL400/XL500 disk shelves.



## Syslog Message

Improper disk drive placement in channel %s disk shelf %d. Loop integrity requires certain disk shelf bays to be populated.

## Parameters

**ChannelName** (STRING): Name of the affected adapter port or switch port loop.

**ShelfID** (INT): ID of the disk shelf.

## ses.drive.missingFromLoopMap

### Severity

ALERT

### Description

This message occurs when a disk shelf detects a disk but the storage adapter cannot detect the disk. (Note: If the message "ses.drive.possShelfAddr" subsequently appears, follow the corrective actions in that message.)

### Corrective Action

Monitor the storage system for approximately 15 minutes. If a "monitor.shelf.accessError.ok" event appears, then the condition has cleared, and no corrective action is needed. If the problem occurs to multiple disks on the same loop or is not resolved in 15 minutes, check the I/O modules at the back of the disk shelves on that loop for errors.

## Syslog Message

The following device attached through adapter %s have not taken expected addresses: %s

## Parameters

**ChannelName** (STRING): Channel the disk shelf is on.

**DiskList** (STRING): List of disks missing from Loop Map.

## ses.drive.missingFromSESMap

### Severity

ALERT

### Description

This message occurs when the storage adapter detects a disk on the Fibre Channel loop or the SAS domain, but enclosure services in the disk shelf does not report a disk with that address. This error indicates that the disk took an address other than what the disk shelf should have provided, or that enclosure services in the disk shelf cannot be contacted for address information. (Note: If the message "ses.drive.possShelfAddr" subsequently appears, follow the corrective actions in that message.) In this condition, the SCSI Enclosure Services (SES) process in the system might be unable to perform certain operations on the disk, such as setting failure LEDs or detecting disk swaps.

### Corrective Action

If the disk is not part of an active aggregate or volume, pull out the disk and push it back into the slot to ensure that the disk is in place. If the problem persists, insert a spare disk into the slot. If the problem is then gone, replace the original disk with the spare one. If the problem occurs to multiple disks on the same loop, check the I/O modules at the back of the disk shelves on that loop for errors. SAS domain disk shelves, DS14mk2 AT disk shelves, and SAS/SATA disk shelves do not rely on disk placement for enclosure services, since they have dedicated SES devices.

## Syslog Message

On adapter %s, the device(s) with the following IDs are detected but are not reported as present by the shelf which should have supplied the IDs: %s

## Parameters

**ChannelName** (STRING): Channel the disk shelf is on.

**LoopIds** (STRING): Loop or SCSI/SAS IDs of the missing disks.

## ses.drive.PossShelfAddr

### Severity

ERROR

### Description

This message occurs in conjunction with the messages ses.drive.missingFromSESmap or ses.drive.missingFromLoopMap when there are devices that have apparently taken a wrong address; the adapter shows device addresses that SCSI Enclosure Services (SES) indicates should not exist, and vice versa. This error is not a fatal condition. It means that SES cannot perform certain operations on the affected disk drives, such as setting failure LEDs, because it is not certain which disk shelf the affected disk drive is in.

### Corrective Action

If the problem is throughout the disk shelf, replace the disk shelf. If the error is only one disk drive per disk shelf, the drive might have taken an incorrect address at power-on. Arrange to make this disk drive a spare, and then reseal it to cause it to take its address again. If the problem persists, insert a different spare disk drive into the slot. If the error then clears, replace the original disk drive. If the problem still persists, there is a hardware problem with the individual disk bay. Replace the disk shelf.

## Syslog Message

Possible actual locations for devices with unexpected addresses: %s.

## Parameters

**Msg** (STRING): Physical locations of disks that do not have expected addresses in the adapter's map of disks.

## ses.fw events

### ses.fw.emptyFile

### Severity

ERROR

### Description

This message occurs when a firmware file is found to be empty during a disk shelf firmware update.

### Corrective Action

Obtain the correct firmware file from NetApp and place it in the /etc/shelf\_fw directory. You can download the firmware file from the NetApp Support Site at <http://support.netapp.com>.

## Syslog Message

Shelf firmware file %s was empty.

## Parameters

**fileName** (STRING): Name of the empty file.

## ses.fw.resourceNotAvailable

### Severity

ERROR

### Description

This message occurs when there is not enough contiguous memory available to download disk shelf firmware.

### Corrective Action

Reduce the amount of system activities before performing a manual disk shelf firmware update. If this does not resolve the issue, reboot the system.

## Syslog Message

Not enough memory available to download disk shelf firmware %s.

## Parameters

**fileName** (STRING): Name of the firmware file.

## ses.giveback events

### ses.giveback.restartAfter

#### Severity

INFORMATIONAL

#### Description

This message occurs when SCSI Enclosure Services (SES) is restarted after giveback.

#### Corrective Action

(None).

#### Syslog Message

Enclosure Services restarting after release of reservations.

#### Parameters

(None).

### ses.giveback.wait

#### Severity

INFORMATIONAL

## Description

This message occurs when SCSI Enclosure Services (SES) information is not available because the system is waiting for giveback.

## Corrective Action

(None).

## Syslog Message

Enclosure Services will be unavailable while waiting for giveback.

## Parameters

(None).

# ses.inconsistent events

## ses.inconsistent.config

### Severity

ALERT

### Description

This message occurs when the system detects an inconsistent configuration of shelf-to-channel connections. One or more shelves on the reported channel are connected to a different set of channels (adapter ports or switch ports) than other shelves in the stack.

### Corrective Action

Verify that all cables are fully inserted on the storage shelves connected to the channel. Verify that the shelves are cabled so that all shelves in a stack are connected to the same channels.

### Syslog Message

Inconsistent configuration detected for shelf connected to adapter %s.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port where shelf connection inconsistencies are found.

# ses.loop events

## ses.loop.configErr

### Severity

ALERT

### Description

This message occurs when the system finds a disk shelf with a Loop Resiliency Circuit (LRC) module before a disk shelf with an Electronically Switched Hub (ESH) module in the loop map.

### Corrective Action

Reorder the disk shelves so that the LRC disk shelf is found after the ESH disk shelf.

## Syslog Message

Unsupported loop configuration for disk shelves attached to channel %s: A %s (ID %d) was found ahead of a %s (ID %d).

## Parameters

**channelName** (STRING): Name of the channel the disk shelves are attached to.

**prodChannel1** (STRING): Disk shelf identity (channel, product, and disk shelf number).

**shelfId1** (INT): Disk shelf ID of the disk shelf on the unsupported loop.

**prodChannel2** (STRING): Disk shelf identity (channel, product, and disk shelf number).

**shelfId2** (INT): Disk shelf ID of the disk shelf on the unsupported loop.

## ses.loop.configOk

### Severity

INFORMATIONAL

### Description

This message occurs when an unsupported Fibre Channel loop configuration error is corrected.

### Corrective Action

(None).

## Syslog Message

Unsupported loop configuration for disk shelves attached to channel %s was corrected.

## Parameters

**channelName** (STRING): Name of the channel on which the unsupported loop configuration error was corrected.

## ses.loop.loopMapUnavailable

### Severity

INFORMATIONAL

### Description

This message occurs when ONTAP® software cannot contact the disk shelf module on a disk shelf and, therefore, cannot resolve the loop map for a disk shelf.

### Corrective Action

For Fibre Channel disk shelves, check whether disk drives are connected to slots 0 and 1 of the disk shelf module and that the disk drives are functional. If not, replace the disk drives. For all other disk shelf types, or if the preceding actions do not correct the error for Fiber Channel disk shelves, reseal the disk shelf module and check if the error is corrected. If not, replace the shelf module.

## Syslog Message

Loop map unavailable for loop on channel %s. Could not check for unsupported disk shelf configurations.

## Parameters

**channelName** (STRING): Name of the channel on which the loop map was unavailable.

# ses.mismatch events

## ses.mismatch.fw.version

### Severity

ERROR

### Description

This message occurs when the two disk shelf modules within the same disk shelf are running different versions of the firmware.

### Corrective Action

Download the latest disk shelf firmware to the disk shelf module that is not up-to-date.

### Syslog Message

The disk shelf modules on disk shelf %s.%d are running two different firmware versions. Disk shelf module A is running %s, and disk shelf module B is running %s.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port.

**shelfID** (INT): ID of the disk shelf.

**revision\_str** (STRING): Revision number for disk shelf module A, the top disk shelf module.

**revision\_str2** (STRING): Revision number for disk shelf module B, the bottom disk shelf module.

# ses.multipath events

## ses.multipath.notSupported

### Severity

NOTICE

### Description

This message occurs when the AT-FCX disk shelf module does not support the multipath high-availability (Multipath HA) configuration.

### Corrective Action

Replace this disk shelf module with one that supports the high-availability multipath configuration.

### Syslog Message

AT-FCX disk shelf module %s on disk shelf %d on channel %s does not support the-high availability multipath configuration. Replace the module with one that supports the-high availability multipath configuration.

### Parameters

**shelfSlot** (STRING): Slot position the disk shelf module is in, either A for the top slot or B for the bottom slot.

**shelfId** (INT): Disk shelf identification number of the disk shelf.

**channelName** (STRING): Name of the channel of the disk shelf.

## ses.multipath.ReqError

### Severity

ALERT

### Description

This message occurs when the system detects that a SAS disk shelf is in a single-path cabling configuration.

### Corrective Action

Check for SAS disk shelf enclosures with only a single path by using the 'sasadmin expander\_map' command. Physically inspect all SAS cables on the attached storage for secure and correct connection.

### Syslog Message

SAS disk shelf detected without a multipath configuration.

### Parameters

(None).

## ses.psu events

### ses.psu.comboError

### Severity

ALERT

### Description

This message occurs when the power supplies that are installed in a disk shelf are not compatible with each other. Prolonged use of incompatible power supplies can result in failures of power supply units (PSUs).

### Corrective Action

Replace one PSU immediately with a compatible model. Use the "storage shelf show -psu-part-number" command to locate all power supplies on the disk shelf, using the part numbers provided in the message as the required command parameter. Consult the NetApp Hardware Universe for compatible part number information.

### Syslog Message

Incompatible power supplies were found on disk shelf "%d", serial number: %s, product ID: %s, PSU part numbers: %s, %s.

### Parameters

**ShelfId** (INT): ID of the disk shelf.

**serialNumber** (STRING): Serial number of the storage shelf enclosure.

**ProdId** (STRING): Product identification of the shelf.

**partNumber1** (STRING): Part number of PSU 1 on the shelf.

**partNumber2** (STRING): Part number of PSU 2 on the shelf.

### ses.psu.coolingReqError

## Severity

ALERT

## Description

This message occurs when the installed power supplies are placed so that air flow requirements of the disk shelf are not met. Placement of power supplies and their integrated fans is part of the disk shelf cooling and air flow design.

## Corrective Action

Install power supplies in the indicated slots, cable to a power source, and turn them on. If the `ses.psu.powerReqError` is not presently posted for the system, an existing supply in the shelf can be moved to fill the required slot. A new instance of the error can be ignored while the supply is moved, but it should clear within a few minutes. If it does not, the condition must be corrected.

## Syslog Message

Improper power supply placement in channel %s disk shelf %d. Power supplies must be present in power supply bays %s to satisfy air flow requirements.

## Parameters

**ChannelName** (STRING): Name of the affected adapter port or switch port.

**ShelfID** (INT): ID of the disk shelf.

**CoolBays** (STRING): PSU bays where PSUs must be installed.

## ses.psu.powerReqError

## Severity

EMERGENCY

## Description

This message occurs when too few power supplies are installed to redundantly satisfy the current-draw requirements of the shelf modules and disk drives in the shelf configuration. The number and type of disk modules and disk drives installed affect the requirements.

## Corrective Action

Install the required number of power supplies in the disk shelf. Ensure that they are fully inserted, cabled to a power source, and turned on.

## Syslog Message

Not enough power supplies are present in channel %s disk shelf %d to satisfy disk drive and shelf power requirements. The shelf requires %d supplies, but only %d are present.

## Parameters

**ChannelName** (STRING): Name of the affected adapter port or switch port.

**ShelfID** (INT): ID of the disk shelf.

**NumReqd** (INT): Number of power supplies required.

**NumPresent** (INT): Number of PSUs present.

## ses.shelf events



## ses.shelf.bay.populate

### Severity

INFORMATIONAL

### Description

This message occurs when one of the drive in bay 1 or 2 is missing.

### Corrective Action

(None).

### Syslog Message

Disk shelf %s is recommended to have slot 1 and slot 2 populated at all times.

### Parameters

**shelfName** (STRING): Disk shelf name.

## ses.shelf.changedId

### Severity

ERROR

### Description

This message occurs on a SAS disk shelf when the disk shelf ID changed after power is applied to the disk shelf.

### Corrective Action

Verify that the disk shelf ID displayed in this message is the same as the disk shelf ID shown on the disk shelf. If they are different and the disk shelf ID displayed in this message is the one you want, reset the disk shelf number to match it. If you want the new disk shelf ID instead of the disk shelf ID displayed in the message, verify that the disk shelf ID you want does not conflict with other disk shelves in the domain. Then power-cycle the disk shelf chassis. You can wait to perform this procedure until your next maintenance window. If the warning persists on both disk shelf modules after you completed the procedure, replace the disk shelf chassis. If it persists on only one disk shelf module, replace the disk shelf module.

### Syslog Message

Disk shelf module %d on disk shelf %s.%d indicates that the disk shelf ID has been changed since the disk shelf module was powered on. %s.

### Parameters

**moduleNumber** (INT): Element number of this module: either 1 for Module A or 2 for Module B.

**channelName** (STRING): Name of the adapter port or switch port associated with the disk shelf; for example, "7a" or "myswitch:5".

**shelfNumber** (INT): Disk shelf ID reported by the disk shelf before the ID change on the thumbwheel switch.

**locationText** (STRING): Location of the disk shelf module on the disk shelf.

## ses.shelf.ctrlFailErr

### Severity

ALERT

## Description

This message occurs when SCSI Enclosure Services' (SES) control to a disk shelf fails.

## Corrective Action

Check the LEDs on the disk shelf and the disk shelf modules on the back of the disk shelf to see whether there are any abnormalities. If the modules appear to be problematic, replace the applicable module. If the SES target is a disk drive, check to see whether the disk drive failed. If it failed, replace the disk drive.

## Syslog Message

Enclosure control failed via %s (disk shelf %d). Alternate path will be used if available.

## Parameters

**SESTarget** (STRING): Address of the SES target for which the SES control failed.  
**shelfID** (INT): Disk shelf ID of the disk shelf that failed.

## ses.shelf.disable.block.request

### Severity

ALERT

### Description

This message is generated when there is a failure while disabling the auto power cycle feature for all shelves that support auto power cycle feature on the system.

### Corrective Action

(None).

### Syslog Message

Fail while disabling shelf auto power cycle capability during %s (%d).

### Parameters

**requestOwner** (STRING): Name of the requestor who requests to disable the auto power cycle feature.  
**status** (INT): Failure status.

## ses.shelf.drawer.min.drives

### Severity

INFORMATIONAL

### Description

This message occurs when one of the first 4 drives are missing in a drawer.

### Corrective Action

Please insert a disk in one of the first 4 bays of the indicated drawer.

### Syslog Message

Disk shelf %d, drawer %d, bay %d missing, drawer needs first 4 drives installed.

### Parameters

**shelfID** (INT): Disk shelf ID.  
**drawerNum** (INT): Disk drawer number.

**bayNum** (INT): Disk bay number.

## **ses.shelf.drawer.open**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when one of the drawers of a disk shelf has been left open.

### **Corrective Action**

Please close the drawer of the indicated drawer. If a drawer has been left open for an extended period of time, this can affect the cooling of components within the disk shelf.

### **Syslog Message**

Disk shelf %d has drawer %d open.

### **Parameters**

**shelfID** (INT): Disk shelf ID.

**drawerNum** (INT): Disk drawer number.

## **ses.shelf.drwr.phy.rateunkwn**

### **Severity**

ALERT

### **Description**

This message occurs when one of the PHY ports of a shelf drawer is reporting an unknown link rate. The PHY port is a physical port in the drawer of the storage shelf that resides on the shelf side. This message indicates an internal error might have occurred in the shelf firmware or software. Communication between the PHY ports and the disk drives in the shelf might have failed or be in degraded states.

### **Corrective Action**

Contact NetApp technical support for assistance.

### **Syslog Message**

Drive shelf "%d", module "%s", drawer "%d", PHY "%d" is reporting an unknown link rate.

### **Parameters**

**shelfID** (INT): Drive shelf ID.

**module** (STRING): Drive shelf module.

**drawerNum** (INT): Shelf drawer number.

**phyNum** (INT): Shelf drawer PHY number.

## **ses.shelf.em.ctrlFailErr**

### **Severity**

EMERGENCY

### **Description**

This message occurs when SCSI Enclosure Services' (SES) control to the internal disk drives of a system

fails.

### **Corrective Action**

Run the "storage shelf show" command to see whether that disk shelf is still being actively monitored. If the command indicates a failure, there is a hardware failure in the system's internal disk shelf.

### **Syslog Message**

Enclosure control failed via embedded disk shelf %d.

### **Parameters**

**shelfID** (INT): Disk shelf ID of the embedded disk shelf that failed.

## **ses.shelf.fw.disrUpdFW**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a disk shelf firmware update operation causes temporary but noticeable interruption of disk I/O. Current disk shelf firmware levels do not support nondisruptive upgrade (NDU).

### **Corrective Action**

(None).

### **Syslog Message**

SCSI Enclosure Services (SES) disk shelf firmware update disrupts disk I/O. The current shelf firmware (product revision level %4s) does not support nondisruptive upgrade (NDU).

### **Parameters**

**ProductRev** (STRING): Product revision level.

## **ses.shelf.fw.disrUpdPath**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a disk shelf firmware update operation causes temporary but noticeable interruption of disk I/O. Redundant paths to the disk shelf are required for nondisruptive upgrade (NDU).

### **Corrective Action**

(None).

### **Syslog Message**

SCSI Enclosure Services (SES) disk shelf firmware update disrupts disk I/O. Redundant paths to the disk shelf are needed for nondisruptive upgrade (NDU) to be used.

### **Parameters**

(None).

## **ses.shelf.fw.update.ndu**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a shelf firmware update is performed using the nondisruptive upgrade (NDU) procedure.

### **Corrective Action**

(None).

### **Syslog Message**

SCSI Enclosure Services (SES) shelf firmware update will be performed using nondisruptive upgrade (NDU).

### **Parameters**

(None).

## **ses.shelf.fw.update.pfu.nrdnt**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a shelf firmware update is performed using the Package Firmware Update (PFU) process because redundant paths to the shelf are not present.

### **Corrective Action**

(None).

### **Syslog Message**

SCSI Enclosure Services (SES) shelf firmware update will be performed using Package Firmware Update (PFU) because redundant paths are not present.

### **Parameters**

(None).

## **ses.shelf.IdBasedAddr**

### **Severity**

ERROR

### **Description**

This message occurs on a serial attached SCSI (SAS) disk shelf when the SAS address of the devices are based on the disk shelf ID instead of the disk shelf backplane serial number. This indicates problems communicating with the disk shelf backplane.

### **Corrective Action**

Reseat the master disk shelf module as indicated by the output of the "storage shelf show" command. If the problem persists, reseat the slave disk shelf module. If the problem persists, find the new master disk shelf module, and replace it. If the problem persists, replace the other disk shelf module. If the problem persists,

replace the disk shelf enclosure.

### Syslog Message

Disk shelf module %d on disk shelf %s.%d indicates that its SAS addresses are based on the disk shelf ID, instead of on the disk shelf backplane serial number. %s.

### Parameters

**moduleNumber** (INT): Element number of this module: either 1 for Module A or 2 for Module B.  
**channelName** (STRING): Name of the adapter port or switch port associated with the disk shelf; for example, "7a" or "myswitch:5".  
**shelfNumber** (INT): Disk shelf ID.  
**locationText** (STRING): Location of the disk shelf module on the disk shelf.

## ses.shelf.invalNum

### Severity

ERROR

### Description

This message occurs when ONTAP® software detects that an SAS disk shelf connected to the system has an invalid disk shelf number.

### Corrective Action

Power-cycle the disk shelf. If the problem persists, replace the disk shelf modules. If the problem still persists, replace the disk shelf.

### Syslog Message

A disk shelf on adapter %s has an invalid disk shelf number of %d. %s.

### Parameters

**channelName** (STRING): Name of the channel the disk shelf is on that is reporting an error.  
**shelfNumber** (INT): Invalid Disk shelf ID.  
**softId** (STRING): String indicating whether a soft disk shelf number is assigned to the disk shelf.

## ses.shelf.maxExceeded

### Severity

ERROR

### Description

This message occurs when more disk shelves are attached than the controller can monitor. The controller cannot detect or identify abnormal conditions in excess shelves, nor can it download firmware to unmonitored shelves. The reported device address is one instance of a device in an unmonitored disk shelf. There might be multiple unmonitored shelves.

### Corrective Action

Consult system documentation for the maximum number of supported disk shelves, which might be lower than the number reported here. Reduce the number of attached disk shelves to no more than the documented number. Consolidating disk drives into sparsely populated shelves and removing excess shelves is one way to reduce shelf count.

## Syslog Message

More than %d disk shelves are attached. Shelf containing %s and possibly other shelves are not monitored.

## Parameters

**max\_enclosures** (INT): Maximum number of disk shelves this system can monitor and control.  
**device** (STRING): Address of the device.

## ses.shelf.mmErr

### Severity

ALERT

### Description

This message occurs when a disk shelf is not supported by the platform it was booted on.

### Corrective Action

Check whether the current version of ONTAP® software supports the disk shelf. If not, install an ONTAP version that supports the disk shelf. If the disk shelf is supported, the error might be cleared by hourly attempts by ONTAP to establish proper contact with the disk shelf.

## Syslog Message

SCSI Enclosure Services (SES) device accessed by %s (product revision level [%s]) (enclosure vendor ID [%s]) (enclosure product ID [%s]) is not supported on this platform.

## Parameters

**DeviceName** (STRING): Address of the unqualified Enclosure Services device.  
**ProductRev** (STRING): Product revision level.  
**EnclosureVendorID** (STRING): Enclosure Vendor Identification.  
**EnclosureProductID** (STRING): Enclosure Product Identification.

## ses.shelf.OSmmErr

### Severity

ALERT

### Description

This message occurs when there are incompatible ONTAP® software versions in a Shared Storage configuration that would cause SCSI Enclosure Services (SES) not to function properly.

### Corrective Action

Update the system that has an earlier ONTAP version to match the one that has the latest ONTAP version.

## Syslog Message

Incompatible versions of ONTAP are running on the different systems in a Shared Storage Configuration.

## Parameters

(None).

## ses.shelf.powercycle.done

**Severity**

INFORMATIONAL

**Description**

This message occurs when a disk shelf power-cycle finishes.

**Corrective Action**

(None).

**Syslog Message**

Disk shelf power-cycled; resuming disk I/O.

**Parameters**

(None).

**ses.shelf.powercycle.start****Severity**

INFORMATIONAL

**Description**

This message occurs when a disk shelf is power-cycled and SCSI Enclosure Services (SES) needs to wait for it to finish.

**Corrective Action**

(None).

**Syslog Message**

Waiting for disk shelf power-cycle to be complete.

**Parameters**

(None).

**ses.shelf.RAS.suggestion****Severity**

INFORMATIONAL

**Description**

This event is generated when various different supported combinations of ESHx type storage controller I/O modules are used in the various different ESHx type storage enclosures. Though this could be a supported configuration, to make the best of the power cycle features it is recommended to keep the FC loops as homogenous as possible.

**Corrective Action**

For best practices, use only the ESH4 controller modules in the DS14-Mk4-FC enclosures.

**Syslog Message**

ses.shelf.RAS.suggestion INFO Channel %s, Mixing %s I/O modules is fully supported on a FC loop. However, for improved reliability migration to homogeneous ESH4 configurations is recommended.



## Parameters

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf.  
**supportedControllerModules** (STRING): ESHx enclosure modules, ESH/ESH2/ESH4

## ses.shelf.reset.auto.fail

### Severity

ERROR

### Description

This event is generated when the automated power cycle attempt in the shelf enclosure - identified by the channel.shelf\_id - had failed. The trigger condition for the automated reset is also reported.

### Corrective Action

Try the 'storage power\_cycle shelf' on the same shelf. If it still fails, please refer to the "Trouble Shooting Section" of the NDR documentation.

### Syslog Message

%s disk shelf %s.%d [%s] failed to perform an automatic reset (0x%x)

## Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.  
**channelName** (STRING): Name of the adapter port or switch port of the disk shelf that was attempted to be power cycled.  
**shelfId** (INT): ID of the disk shelf that was attempted to be power cycled.  
**serialNumber** (STRING): Serial number of the storage shelf enclosure.  
**triggerCondition** (INT): Trigger reason for the automated reset, as reported by the shelf enclosure.

## ses.shelf.reset.auto.good

### Severity

INFORMATIONAL

### Description

This event is generated when the automated power cycle attempt in the shelf enclosure - identified by the channel.shelf\_id - had succeeded. The trigger for the automated reset is also reported.

### Corrective Action

(None).

### Syslog Message

%s disk shelf %s.%d [%s] has recovered from an automatic reset (0x%x)

## Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.  
**channelName** (STRING): Name of the adapter port or switch port of the disk shelf that was power cycled.  
**shelfId** (INT): ID of the disk shelf that was power cycled.  
**serialNumber** (STRING): Serial number of the storage shelf enclosure.  
**triggerCondition** (INT): Trigger reason for the automated reset, as reported by the shelf enclosure.

## ses.shelf.reset.capability

### Severity

INFORMATIONAL

### Description

This event is generated when the periodic storage enclosure services monitor detects some of the DS14-Mk4-FC, DS14-Mk2-FC and DS14-Mk2-AT storage enclosures are not being upgraded to the correct requirements, to avail the enhanced power cycle features. This is just an informational message only, not an error. The idea is to make the admin follow the suggested best practices, so that the Filer could make use of the advanced features provided.

### Corrective Action

Check the power supply unit information on all the DS14-Mk4-FC, DS14-Mk2-FC and the DS14-Mk2-AT shelf enclosure modules by executing the 'environment shelf' command. Plan to replace the power supply units that are not "Type:HRE capable".

### Syslog Message

Some %s disk shelves on this system are capable of automatic power reset, but some do not have appropriate firmware or PSU model. For improved system RAS, consider upgrading shelves to have this capability. The environ shelf command shows this feature's status.

### Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.

## ses.shelf.reset.request

### Severity

INFORMATIONAL

### Description

This event is generated when the command 'storage power\_cycle shelf start' is successfully sent to the shelf - identified by the channel.shelf\_id, and is of type DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.

### Corrective Action

(None).

### Syslog Message

Reset requested on %s disk shelf %s.%d [%s].

### Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf to be power cycled.

**shelfId** (INT): ID of the disk shelf that is to be power cycled.

**serialNumber** (STRING): Serial number of the storage shelf enclosure.

## ses.shelf.reset.request.fail

### Severity

ERROR

## Description

This event is generated when the command 'storage power\_cycle shelf start' sent to the shelf enclosure - identified by the channel.shelf\_id - could not power cycle the shelf enclosure successfully.

## Corrective Action

Try the 'storage power\_cycle shelf start' on the same shelf again. If it still fails, please refer to the "Trouble Shooting Section" of the NDR documentation.

## Syslog Message

%s disk shelf %s.%d [%s] failed to reset when requested.

## Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf that was attempted to be power cycled.

**shelfId** (INT): ID of the disk shelf that was attempted to be power cycled.

**serialNumber** (STRING): Serial number of the storage shelf enclosure.

## ses.shelf.reset.request.good

### Severity

INFORMATIONAL

### Description

This event is generated when the command 'storage power\_cycle shelf start' sent to the shelf - identified by the channel.shelf\_id, and is of type DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT - has successfully recovered.

### Corrective Action

(None).

### Syslog Message

%s disk shelf %s.%d [%s] has recovered from a requested reset.

### Parameters

**shelfType** (STRING): Shelf enclosure type, DS14-Mk4-FC, DS14-Mk2-FC or DS14-Mk2-AT.

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf that was power cycled.

**shelfId** (INT): ID of the disk shelf that was power cycled.

**serialNumber** (STRING): Serial number of the storage shelf enclosure.

## ses.shelf.reset.unsup.drives

### Severity

ERROR

### Description

This message occurs when Nondisruptive Recovery (NDR) functionality is disabled due to the presence of Parallel ATA (PATA) disk drives on the disk loop.

### Corrective Action

Replace all X266\_MTOMC320PSX or X262\_MCALY250PSX disk drives and then turn on the automatic reset capability.

## Syslog Message

NDR functionality is disabled due to the presence of X266\_MTOMC320PSX or X262\_MCALY250PSX disk drives on the disk loop.

## Parameters

(None).

## ses.shelf.sameNumReassign

### Severity

ERROR

### Description

This message occurs when ONTAP® software detects more than one SAS shelf connected to the same adapter has the same disk shelf number.

### Corrective Action

Change the disk shelf number on the disk shelf to one that does not conflict with other disk shelves attached to the same adapter. Then halt the system and reboot the disk shelf.

## Syslog Message

Multiple disk shelves with disk shelf number %d are found on channel %s. Disk shelf with serial number %s has been assigned the soft disk shelf number of %d.

## Parameters

**prevId** (INT): Conflicting disk shelf number.

**channelName** (STRING): Name of the channel the disk shelf is on that is reporting an error.

**serialNum** (STRING): Serial number of the disk shelf that is being assigned the soft ID.

**softId** (INT): Newly assigned soft disk shelf number.

## ses.shelf.unsupportAllowErr

### Severity

ALERT

### Description

This message occurs when a disk shelf is not supported by ONTAP® software. ONTAP will continue to use the disk shelf, but environmental monitoring of the disk shelf is not possible.

### Corrective Action

Check whether the current version of ONTAP supports this disk shelf. If not, install an appropriate version. If ONTAP does support this disk shelf, this error might be cleared on hourly attempts by ONTAP to establish proper contact with the disk shelf.

## Syslog Message

Unsupported disk shelf found on channel %s.

## Parameters

**ChannelName** (STRING): Name of the channel that contains the disk shelf.

## ses.shelf.unsupportedErr

### Severity

ALERT

### Description

This message occurs when a disk shelf is not supported by ONTAP® software.

### Corrective Action

Check whether the current version of ONTAP supports this disk shelf. If not, install an appropriate version.

### Syslog Message

Unsupported disk shelf found. Taking channel %s offline.

### Parameters

**ChannelName** (STRING): Name of the channel that contains the disk shelf. This channel will be taken offline.

## ses.status.events

### ses.status.ACPErr

### Severity

ALERT

### Description

This message occurs when a failure has been detected in ACP processor that provides alternate control path capability.

### Corrective Action

Replace the module. In some disk shelf types, this function is integrated into the SAS interface modules.

### Syslog Message

%s ACP Processor error for %s %d: %s %s %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

### ses.status.ACPIInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a problem previously reported about the ACP Processor has been corrected or when other information about the ACP Processor that does not necessarily require customer action is

reported.

### Corrective Action

(None).

### Syslog Message

%s ACP Processor information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**infoMsg** (STRING): Generic SES information.

## ses.status.ACPWarn

### Severity

ERROR

### Description

This message occurs when a non-fatal condition has been detected in the ACP Processor that provides alternate control path capability.

### Corrective Action

Replace the module. In some disk shelf types, this function is integrated into the SAS interface modules.

### Syslog Message

%s ACP Processor warning for %s %d: %s. %s %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**warningMsg** (STRING): Generic SES Warning.

**warningText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.ATFCXError

### Severity

ALERT

### Description

This message occurs when the reporting disk shelf detects an error in the indicated AT-FCX module. The module might not be able to perform I/O to disks within the disk shelf.

### Corrective Action

Verify that the AT-FCX module is fully seated and secured. If the problem persists, replace the AT-FCX module.

## Syslog Message

%s error for %s SES controller module %s: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).

**typeText** (STRING): SES type descriptor text for this AT-FCX module.

**module** (STRING): Module A is the upper module in the enclosure on the back side. Module B is the lower module in the enclosure on the back side.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Description of the exact warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.ATFCXInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a previously reported error in the AT-FCX module is corrected, or the system reports other information that does not necessarily require customer action.

### Corrective Action

(None).

## Syslog Message

%s information for %s SES controller module %s: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**module** (STRING): Module A is the upper module in the enclosure on the back side. Module B is the lower module in the enclosure on the back side.

**infoMsg** (STRING): Generic SES information.

## ses.status.battery.error

### Severity

ERROR

### Description

This message occurs when a storage shelf reports a generic "battery error detected" error. The error could be due to an inability to read the A/D convertor or because the intermodule connections are down.

### Corrective Action

If the intermodule connection is down on one or both of the PCIe Switch Modules (PSMs), bring the intermodule connection up (by rebooting the affected module). If the intermodule connection is up on both of the PSMs (there are no errors on both PSMs), replace the coin cell battery. If the issue persists, a module might need to be replaced. Contact NetApp technical support.

## Syslog Message

%s battery failure error for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the battery.  
**batteryNumber** (INT): SES element number of the battery.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the error.  
**locationText** (STRING): Location of the element on the storage shelf.

## ses.status.battery.info

### Severity

NOTICE

### Description

This message occurs when a condition that was previously reported about a storage shelf battery has been corrected or when other information about the battery that might not require customer action is reported.

### Corrective Action

(None).

### Syslog Message

%s battery information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the battery.  
**batteryNumber** (INT): SES element number of the battery.  
**infoMsg** (STRING): Generic SES information.

## ses.status.battery.missing

### Severity

ERROR

### Description

This message occurs when a storage shelf reports that a battery is not installed.

### Corrective Action

Install the battery. If the problem persists, contact NetApp technical support.

### Syslog Message

%s battery warning for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the battery.  
**batteryNumber** (INT): SES element number of the battery.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the battery on the storage shelf.



## ses.status.battery.threshErr

### Severity

ALERT

### Description

This message occurs when the element threshold of a battery on a storage shelf is outside the threshold limits.

### Corrective Action

Install a new battery. If the problem persists, this indicate a hardware failure. Contact NetApp technical support.

### Syslog Message

%s battery failure error for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the battery.

**batteryNumber** (INT): SES element number of the failing battery.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact description of the error.

**locationText** (STRING): Location of the battery on the storage shelf.

## ses.status.bootDv.error

### Severity

ERROR

### Description

This message occurs when a storage shelf reports a boot device SMART error on a PCIe Switch Module (PSM)

### Corrective Action

Replace the boot device. If the issue persists, contact NetApp technical support.

### Syslog Message

%s boot device failure for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the boot device.

**bootDvNumber** (INT): SES element number of the boot device.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Description of the error.

**locationText** (STRING): Location of the boot device on the storage shelf.

## ses.status.bootDv.info

### Severity

NOTICE

## Description

This message occurs when a condition that was previously reported about a storage shelf boot device has been corrected or when other information about the boot device that might not require customer action is reported.

## Corrective Action

(None).

## Syslog Message

%s boot device notification for %s %d: %s.

## Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the boot device.

**bootDvNumber** (INT): SES element number of the boot device.

**infoMsg** (STRING): Generic SES information.

## ses.status.bootDv.missing

### Severity

ERROR

### Description

This message occurs when the PCIe Switch Module (PSM) of a storage shelf reports that a partner PSM boot device is not installed.

### Corrective Action

Install the boot device on the partner PSM. If the problem persists, contact NetApp technical support.

### Syslog Message

%s partner %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the boot device.

**bootDvNumber** (INT): SES element number of the boot device.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Description of the error.

**locationText** (STRING): Location of the boot device on the storage shelf.

## ses.status.connectorInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a problem previously reported about the SAS connector has been corrected.

### Corrective Action

(None).

## Syslog Message

%s SAS connector information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**infoMsg** (STRING): Generic SES information.

## ses.status.connectorWarn

### Severity

ERROR

### Description

This message occurs when a non-fatal condition has been detected in communication with the SAS connector or its EEPROM. This error does not necessarily mean that the data path through the connector is faulty.

### Corrective Action

Ensure that the connector is completely installed. If the error persists, replace the cable first, then the disk shelf I/O module.

## Syslog Message

%s SAS connector warning for %s %d: %s. %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**warningMsg** (STRING): Generic SES Warning.  
**warningText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.currentError

### Severity

ERROR

### Description

This message occurs when a critical condition is detected in the indicated disk shelf current sensor. The disk shelf might be able to continue operation.

### Corrective Action

Verify that the PSU and the AC line are supplying power. Monitor the power grid for abnormalities. If all seems to be normal, replace the PSU. If the problem persists, contact NetApp technical support.

## Syslog Message

%s current error for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this power supply.  
**psNumber** (INT): SES element number of the power supply.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the exact warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.currentInfo

### Severity

INFORMATIONAL

### Description

This message occurs when an error or warning condition previously reported by or about the disk shelf current sensor is corrected, or the system reports other information about the current in the disk shelf that does not necessarily require customer action.

### Corrective Action

(None).

### Syslog Message

%s current information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this sensor.  
**sensorNumber** (INT): SES element number of the reporting sensor.  
**infoMsg** (STRING): Generic SES information.

## ses.status.currentWarning

### Severity

ERROR

### Description

This message occurs when a warning condition is detected in the indicated disk shelf current sensor. The disk shelf might be able to continue operation.

### Corrective Action

Verify that the PSU and the AC line are supplying power. Monitor the power grid for abnormalities. If all seems to be normal, replace the PSU. If the problem persists, contact NetApp technical support.

### Syslog Message

%s current warning for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this power supply.  
**psNumber** (INT): SES element number of the power supply.  
**errorMsg** (STRING): generic SES error.  
**errorText** (STRING): Description of the exact warning.

**locationText** (STRING): Location of the module on the shelf.

## **ses.status.dimmed.error**

### **Severity**

ERROR

### **Description**

This message occurs when a storage shelf reports a generic "DIMM error detected" error on a PCIe Switch Module (PSM). The error could be because intermodule connections are down.

### **Corrective Action**

If the intermodule connection is down on one or both of the PCIe Switch Modules (PSMs), bring the intermodule connection up (by rebooting the affected module). If the intermodule connection is up on both of the PSMs (there are no errors on both PSMs), replace the DIMM. If the issue persists, a module might need to be replaced. Contact NetApp technical support.

### **Syslog Message**

%s DIMM failure for %s %d: %s%s. %s.

### **Parameters**

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the DIMM.  
**dimmedNumber** (INT): SES element number of the DIMM.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the error.  
**locationText** (STRING): Location of the DIMM on the storage shelf.

## **ses.status.dimmed.info**

### **Severity**

NOTICE

### **Description**

This message occurs when a condition that was previously reported about a storage shelf DIMM has been corrected or when other information about the DIMM that might not require customer action is reported.

### **Corrective Action**

(None).

### **Syslog Message**

%s DIMM notification for %s %d: %s.

### **Parameters**

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the DIMM.  
**dimmedNumber** (INT): SES element number of the DIMM.  
**infoMsg** (STRING): Generic SES information.

## ses.status.dimmmissing

### Severity

ERROR

### Description

This message occurs when the PCIe Switch Module (PSM) of a storage shelf reports that a DIMM is not installed.

### Corrective Action

Install the DIMM. If the problem persists, contact NetApp technical support.

### Syslog Message

%s %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the DIMM.  
**dimmmNumber** (INT): SES element number of the DIMM.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the error.  
**locationText** (STRING): Location of the DIMM on the storage shelf.

## ses.status.dimmm.threshErr

### Severity

ALERT

### Description

This message occurs when the element threshold of a DIMM is outside the threshold limits in the indicated PCIe Switch Module (PSM) of a storage shelf.

### Corrective Action

Install a new DIMM. If the problem persists, this indicates a hardware failure. Contact NetApp technical support.

### Syslog Message

%s DIMM threshold failure error for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the DIMM.  
**dimmmNumber** (INT): SES element number of the failing DIMM.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the error.  
**locationText** (STRING): Location of the DIMM on the storage shelf.

## ses.status.displayInfo

### Severity

INFORMATIONAL

## Description

This message occurs when a previous condition in the display panel has been corrected.

## Corrective Action

(None).

## Syslog Message

%s display information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this display panel.

**displayNumber** (INT): SES element number of this display panel.

**infoMsg** (STRING): Generic SES information.

## ses.status.displayWarning

### Severity

ERROR

### Description

This message occurs when the enclosure services module has detected a warning condition for the disk shelf display panel. The disk shelf might be unable to provide correct addresses to its disks.

### Corrective Action

If possible, verify that the connection between the disk shelf and the display is secure. Verify that the Enclosure Services module or modules are fully seated; replacing them might solve the problem. If the problem persists, the Enclosure Services module that detected the warning condition might be faulty. If the problem persists after the module or modules are replaced, replace the disk shelf. If the problem persists, contact NetApp technical support for assistance.

### Syslog Message

%s display warning for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this display panel.

**displayNumber** (INT): SES element number of the this display panel.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.driveError

### Severity

ERROR

### Description

This message occurs when a critical condition has been detected for the disk drive in the shelf. The drive might fail.

## Corrective Action

Make sure that the drive is not running on a degraded volume. If it is, then add as many spares as necessary into the system, up to the specified level. After the volume is no longer in degraded mode, replace the drive that is failing.

## Syslog Message

Error on drive %d on %s shelf %s.%d%s.

## Parameters

**driveNumber** (INT): SES element number of the failing drive.  
**shelfType** (STRING): Type of the disk shelf that is experiencing the problem.  
**channel** (STRING): Channel the shelf is on.  
**shelfNumber** (INT): Number of the disk shelf the drive is on.  
**ErrorText** (STRING): Exact error description.

## ses.status.driveOk

### Severity

INFORMATIONAL

### Description

This message occurs when a disk drive that was previously experiencing problem returns to normal operation.

### Corrective Action

(None).

## Syslog Message

The error on drive %d on %s shelf %s.%d has been corrected.

## Parameters

**driveNumber** (INT): SES element number of the drive.  
**shelfType** (STRING): Type of the disk shelf that contains the drive.  
**channel** (STRING): Channel the shelf is on.  
**shelfNumber** (INT): Number of the disk shelf the drive is on.

## ses.status.driveWarning

### Severity

ERROR

### Description

This message occurs when a non-critical condition has been detected for the disk drive in the shelf. The drive might fail.

### Corrective Action

Make sure that the drive is not running on a degraded volume. If it is, then add as many spares as necessary into the system, up to the specified level. After the volume is no longer in degraded mode, replace the drive that is failing.



## Syslog Message

Error on drive %d on %s shelf %s.%d%s.

## Parameters

**driveNumber** (INT): SES element number of the failing drive.  
**shelfType** (STRING): Type of the disk shelf that is experiencing the problem.  
**channel** (STRING): Channel the shelf is on.  
**shelfNumber** (INT): Number of the disk shelf the drive is on.  
**warningText** (STRING): Exact warning description.

## ses.status.electronicsError

### Severity

ALERT

### Description

This message occurs when a failure is detected in a module that provides disk enclosure services monitoring capability.

### Corrective Action

Replace the shelf module.

## Syslog Message

%s environmental monitoring error for %s %d: %s %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.electronicsInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a problem previously reported about the disk shelf enclosure services electronics has been corrected or when other information about the enclosure electronics that does not necessarily require customer action is reported.

### Corrective Action

(None).

## Syslog Message

%s environmental monitoring information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**infoMsg** (STRING): Generic SES information.

## **ses.status.electronicsWarn**

### **Severity**

ERROR

### **Description**

This message occurs when a non-fatal condition is detected in a module that provides disk enclosure services monitoring capability.

### **Corrective Action**

Replace the module. In most disk shelf types, this function is integrated into the FC, SCSI, or SAS interface modules.

### **Syslog Message**

%s environmental monitoring warning for %s %d: %s. %s %s.

### **Parameters**

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**warningMsg** (STRING): Generic SES Warning.

**warningText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## **ses.status.enclError**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a failure is detected in the disk enclosure (shelf) internal hardware.

### **Corrective Action**

Contact NetApp technical support for assistance with shelf or shelf midplane replacement.

### **Syslog Message**

%s environmental monitoring error for %s %d: %s %s %s.

### **Parameters**

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the failed component on the disk shelf.

## ses.status.enclInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a problem previously reported about the disk shelf hardware has been corrected.

### Corrective Action

(None).

### Syslog Message

%s disk inclosure information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**infoMsg** (STRING): Generic SES information.

## ses.status.enclWarn

### Severity

ERROR

### Description

This message occurs when an error condition is present in the disk shelf that causes the shelf fault LED to be lit, including faults in shelf components or disks marked as failed by ONTAP® software.

### Corrective Action

Correct error conditions in the disk shelf. Check that the power supplies are in the proper locations, that there are the required number of them, and that they are powered on. Remove failed disk drives. If multiple shelf modules are expected, install a missing module, and ensure that the firmware revisions are at the same level, preferably the latest revision available. If the error persists, contact NetApp technical support for assistance.

### Syslog Message

%s disk enclosure warning for %s %d: %s. %s %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**warningMsg** (STRING): Generic SES Warning.

**warningText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.etherConn.info

**Severity**

NOTICE

**Description**

This message occurs when a problem previously reported about a PCIe switch module (PSM) port's Ethernet connector has been corrected.

**Corrective Action**

(None).

**Syslog Message**

%s Ethernet connector information for port %s: %s.

**Parameters**

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**connName** (STRING): SES element name of this Ethernet connector.

**infoMsg** (STRING): Generic SES information.

**ses.status.etherConn.warn****Severity**

ERROR

**Description**

This message occurs when a non-fatal condition has been detected in communication with a PCIe switch module (PSM) port's external Ethernet connector or its EEPROM. This error does not necessarily mean that the data path through the connector is faulty.

**Corrective Action**

Verify that the Ethernet cable or transceiver is correctly installed. If the error persists, replace the cable. If necessary, replace the PSM.

**Syslog Message**

%s Ethernet connector warning for port %s: %s. %s %s.

**Parameters**

**prodChannel** (STRING): Storage shelf identity (channel, product, and shelf number).

**connName** (STRING): SES element name of this Ethernet connector.

**warningMsg** (STRING): Generic SES warning.

**warningText** (STRING): Warning text.

**locationText** (STRING): Location of the module on the storage shelf.

**ses.status.fanError****Severity**

EMERGENCY

**Description**

This message occurs when the indicated disk shelf cooling fan or fan module has failed, and the shelf or its components are not receiving required cooling airflow.

## Corrective Action

Verify that the fan module is fully seated and secured (the fan is integrated into the power supply module in some disk shelves). If the problem persists, replace the fan module. If the problem still persists, contact NetApp technical support for assistance.

## Syslog Message

%s cooling fan error for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this fan.

**fanNumber** (INT): SES element number of this fan.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the error.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.fanInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a condition previously reported about the disk shelf cooling fan or fan module has been corrected or when other information about the fans that does not necessarily require customer action is reported.

### Corrective Action

(None).

## Syslog Message

%s cooling fan information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this fan.

**fanNumber** (INT): SES element number of this fan.

**infoMsg** (STRING): Generic SES information.

## ses.status.fanWarning

### Severity

NOTICE

### Description

This message occurs when a disk shelf cooling fan is not operating to specification, or a component of a fan module has stopped functioning. The disk shelf components continue to receive cooling airflow but might eventually reach temperatures that are out of specification.

### Corrective Action

Verify that the fan or fan module is fully seated and secured (fans are integrated into the power supply module in some shelves). If the problem persists, replace the fan module. If the problem still persists, contact NetApp technical support for assistance.

## Syslog Message

%s cooling fan warning for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this fan.

**fanNumber** (INT): SES element number of this fan.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.inputPowerError

### Severity

ALERT

### Description

This message occurs when the system detects a failure in the Input Power Monitor. The Input Power Monitor monitors power to the Power Supply Units (PSU).

### Corrective Action

Check that the PSU is receiving sufficient power from an AC source.

## Syslog Message

%s Input Power Monitor error for %s %d: %s %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the error.

**locationText** (STRING): Location of the PSU on the disk shelf.

## ses.status.inputPowerInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a previously reported Input Power Monitor problem of a disk shelf is corrected, or when the system is reporting information about the Input Power Monitor that does not necessarily require customer action.

### Corrective Action

(None).

## Syslog Message

%s Input Power Monitor information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**infoMsg** (STRING): Generic SES information.

## ses.status.inputPowerWarn

### Severity

ERROR

### Description

This message occurs when the system detects a non-fatal condition in the Input Power Monitor of a disk shelf. The Input Power Monitor monitors power to the Power Supply Units (PSU).

### Corrective Action

Check that the PSU is receiving sufficient power from an AC source.

### Syslog Message

%s Input Power Monitor warning for %s %d: %s. %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**warningMsg** (STRING): Generic SES warning.  
**warningText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the PSU on the disk shelf.

## ses.status.ModuleError

### Severity

ALERT

### Description

This message occurs when the reporting disk shelf has detected an error in the indicated disk shelf module.

### Corrective Action

Verify that the shelf module is fully seated and secure. If the problem persists, replace the disk shelf module.

### Syslog Message

%s %s error for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**moduleType** (STRING): Type of this module.  
**typeText** (STRING): SES type descriptor text for this module.  
**ModuleNumber** (INT): Element number of this module: either 1 for Module A or 2 for Module B.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.ModuleInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a previously reported error in the shelf module has been corrected or when other information that does not necessarily require customer action is reported.

### Corrective Action

(None).

### Syslog Message

%s %s information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**moduleType** (STRING): Type of this module.

**typeText** (STRING): SES type descriptor text for this module.

**ModuleNumber** (INT): Element number of this module: either 1 for Module A or 2 for Module B.

**infoMsg** (STRING): Generic SES information.

## ses.status.ModuleWarn

### Severity

ALERT

### Description

This message occurs when the reporting disk shelf detects a warning condition in the indicated disk shelf module.

### Corrective Action

Verify that the shelf module is fully seated and secure. If the problem persists, replace the disk shelf module.

### Syslog Message

%s %s warning for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**moduleType** (STRING): Type of this module.

**typeText** (STRING): SES type descriptor text for this module.

**moduleNumber** (INT): Element number of this module: either 1 for Module A or 2 for Module B.

**warningMsg** (STRING): Generic SES warning.

**warnText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.powerCrestError

### Severity

ALERT



## Description

This message occurs when the system detects a failure in the Power Crest Factor Monitor for a disk shelf PSU. The Power Crest Factor Monitor will detect issues with the disk shelf power supply unit (PSU).

## Corrective Action

If an abnormal power crest factor is indicated, replace the corresponding PSU.

## Syslog Message

%s Power Crest Factor Monitor error for %s %d: %s %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the PSU on the disk shelf.

## ses.status.powerCrestInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a previously reported Power Crest Factor Monitor problem of a disk shelf is corrected, or when the system is reporting information about the Power Crest Factor Monitor that does not necessarily require customer action.

### Corrective Action

(None).

### Syslog Message

%s Power Crest Factor Monitor information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**infoMsg** (STRING): Generic SES information.

## ses.status.powerCrestWarn

### Severity

ERROR

### Description

This message occurs when the system detects a non-fatal condition in the Power Crest Factor Monitor of a power supply unit (PSU) of a disk shelf.

### Corrective Action

Replace the PSU of the disk shelf.

## Syslog Message

%s Power Crest Factor Monitor warning for %s %d: %s. %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**warningMsg** (STRING): Generic SES warning.

**warningText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the PSU on the disk shelf.

## ses.status.procCplxError

### Severity

ALERT

### Description

This message occurs when the system detects a failure in the Processor Complex. The Processor Complex is the set of electronic hardware, including the CPU, memory, and so on, that runs ONTAP® software.

### Corrective Action

If insufficient power or abnormal voltage is indicated, check the power supply configuration; otherwise, replace the module. In some disk shelf types, this function is integrated into the SAS interface modules.

## Syslog Message

%s Processor Complex error for %s %d: %s %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this element.

**electrNumber** (INT): SES element number of this element.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.procCplxInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a previously reported Processor Complex problem is corrected, or when the system is reporting information about the Processor Complex that does not necessarily require customer action. The Processor Complex is the set of electronic hardware, including the CPU, memory, and so on, that runs ONTAP® software.

### Corrective Action

(None).

## Syslog Message

%s Processor Complex information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**infoMsg** (STRING): Generic SES information.

## ses.status.procCplxWarn

### Severity

ERROR

### Description

This message occurs when the system detects a non-fatal condition in the Processor Complex. The Processor Complex is the set of electronic hardware, including the CPU, memory, and so on, that runs ONTAP® software.

### Corrective Action

Replace the module. In some disk shelf types, this function is integrated into the SAS interface modules.

### Syslog Message

%s Processor Complex warning for %s %d: %s. %s %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this element.  
**electrNumber** (INT): SES element number of this element.  
**warningMsg** (STRING): Generic SES warning.  
**warningText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.psError

### Severity

ALERT

### Description

This message occurs when a critical condition is detected in the indicated storage shelf power supply. The power supply might fail.

### Corrective Action

Verify that input power to the shelf is working correctly. If separate events of this type are reported simultaneously, then the common power distribution point might be at fault. Make sure that the shelf power cords are fully inserted and secured, the supply is fully seated and secured, and the supply is switched on. Verify that power supply fans, if any, are functioning. If the problem persists, replace the power supply. If the problem still persists, contact NetApp technical support.

### Syslog Message

%s power error for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this power supply.

**psNumber** (INT): SES element number of the failing power supply.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Exact description of the error.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.psInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a condition previously reported about the disk shelf power supply has been corrected or when other information about the power supply that does not necessarily require customer action is reported.

### Corrective Action

(None).

### Syslog Message

%s power supply information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this supply.  
**psNumber** (INT): SES element number of the reporting supply.  
**infoMsg** (STRING): Generic SES information.

## ses.status.psWarning

### Severity

ERROR

### Description

This message occurs when a warning condition has been detected in the indicated storage shelf power supply. The power supply might be able to continue operation.

### Corrective Action

Verify that the disk shelf is receiving power. If separate events of this type are reported simultaneously, then the common power distribution point might be at fault. Make sure that the disk shelf power cords are fully inserted and secured, the supply is fully seated and secured, and the supply is switched on. If the problem persists, replace the power supply. If the problem still persists, contact NetApp technical support.

### Syslog Message

%s power warning for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).  
**typeText** (STRING): SES type descriptor text for this supply.  
**psNumber** (INT): SES element number of the erring supply.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Exact wording of the warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.shelfCount.info

### Severity

NOTICE

### Description

This message occurs when the shelf count changes due to shelf removal, insertion or the controller was rebooted.

### Corrective Action

(None).

### Syslog Message

The new shelf count is %d, the old shelf count was %d.

### Parameters

**currentShelfCount** (INT): New shelf count.

**oldShelfCount** (INT): Previous shelf count.

## ses.status.temperatureError

### Severity

ERROR

### Description

This message occurs when the indicated disk shelf temperature sensor reports a temperature that exceeds the specifications for the disk shelf or its components.

### Corrective Action

Verify that the ambient temperature where the shelf is installed is within NetApp equipment specifications using the "storage shelf show" command and that airflow clearances are maintained. If the same disk shelf also reports fan or fan module failures, correct that problem now. If the problem is reported by the ambient temperature sensor (located on the operator panel), verify that the connection between the disk shelf and the panel is secure, if possible. If the problem persists, and if the shelf has multiple temperature sensors of which only one exhibits the problem, replace the module that contains the sensor that reports the error. If problem still persists, contact NetApp technical support for assistance. The temperature thresholds for each shelf can be displayed by using the "storage shelf show" command. command.

### Syslog Message

%s temperature error for %s %d: %s%s. Current temperature: %s C (%s F). %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this sensor.

**thermNumber** (INT): SES element number of this sensor.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the error.

**currentCTemp** (STRING): Temperature reported by the sensor, in Celsius, if available.

**currentFTemp** (STRING): Temperature reported by the sensor, in Fahrenheit, if available.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.temperatureInfo

### Severity

INFORMATIONAL

### Description

This message occurs when an error or warning condition previously reported by or about the disk shelf temperature sensor has been corrected or when other information about the temperature in the disk shelf that does not necessarily require customer action is reported.

### Corrective Action

(None).

### Syslog Message

%s temperature information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this sensor.

**thermNumber** (INT): SES element number of this sensor.

**infoMsg** (STRING): Generic SES information.

## ses.status.temperatureWarning

### Severity

ALERT

### Description

This message occurs when the indicated disk shelf temperature sensor reports a temperature that is close to exceeding the specifications for the disk shelf or its components.

### Corrective Action

Verify that the ambient temperature where the disk shelf is installed is within NetApp equipment specifications by using the "storage shelf show" command, and that airflow clearances are maintained. If this disk shelf also reports fan or fan module errors or warnings, correct those problems now. If the problem persists, and the shelf has multiple temperature sensors of which only one exhibits the problem, replace the module that contains the sensor. If the problem persists, contact NetApp technical support for assistance. The temperature thresholds for each shelf can also be displayed by using the "storage shelf show" command.

### Syslog Message

%s temperature warning for %s %d: %s%s. Current temperature: %s C (%s F). %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this sensor.

**thermNumber** (INT): SES element number of this sensor.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**currentCTemp** (STRING): Temperature reported by the sensor, in Celsius, if available.

**currentFTemp** (STRING): Temperature reported by the sensor, in Fahrenheit, if available.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.upsError

### Severity

ALERT

### Description

This message occurs when the disk shelf has detected a failure in the uninterruptible power supply (UPS) attached to it. This might occur, for example, if power to the UPS has been lost.

### Corrective Action

Restore power to the UPS. Verify that the connection from the UPS to the disk shelf is in place and secured, and that the UPS is enabled. If the problem persists, contact NetApp technical support.

### Syslog Message

%s uninterruptible power error for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this UPS

**upsNumber** (INT): SES element number of the UPS

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the error.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.upsInfo

### Severity

INFORMATIONAL

### Description

This message occurs when a condition previously reported about the uninterruptible power supply (UPS) attached to the disk shelf has been corrected or when other information about the UPS that does not necessarily require customer action is reported.

### Corrective Action

(None).

### Syslog Message

%s uninterruptible power information for %s %d: %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this UPS.

**upsNumber** (INT): SES element number of the reporting UPS.

**infoMsg** (STRING): Generic SES information.

## ses.status.upsWarning

### Severity

ALERT

## Description

This message occurs when the disk shelf has detected a warning condition in the uninterruptible power supply (UPS) attached to it. This might occur, for example, if power to the UPS has been lost.

## Corrective Action

Restore power to the UPS. Verify that the connection from the UPS to the disk shelf is attached and secured, and that the UPS is enabled. If the problem persists, contact NetApp technical support.

## Syslog Message

%s uninterruptible power warning for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and shelf number).

**typeText** (STRING): SES type descriptor text for this UPS.

**upsNumber** (INT): SES element number of the UPS.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Exact wording of the warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.volError

### Severity

EMERGENCY

### Description

This message occurs when a critical condition is detected in the indicated disk shelf voltage sensor. The disk shelf might be able to continue operation.

### Corrective Action

Verify that the power supply unit (PSU) and the AC line are supplying power. Monitor the power grid for abnormalities. If all seems to be normal, replace the PSU. If the problem persists, contact technical support.

### Syslog Message

%s voltage error for %s %d: %s%s. %s.

### Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).

**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this power supply.

**psNumber** (INT): SES element number of the power supply.

**errorMsg** (STRING): Generic SES error.

**errorText** (STRING): Description of the exact warning.

**locationText** (STRING): Location of the module on the disk shelf.

## ses.status.volInfo

### Severity

INFORMATIONAL

### Description

This message occurs when an error or warning condition previously reported by or about the disk shelf voltage sensor is corrected, or the system reports other information about the voltage in the disk shelf that does not necessarily require customer action.



## Corrective Action

(None).

## Syslog Message

%s voltage information for %s %d: %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for this sensor.  
**sensorNumber** (INT): SES element number of the reporting sensor.  
**infoMsg** (STRING): Generic SES information.

## ses.status.volWarning

### Severity

ERROR

### Description

This message occurs when a warning condition is detected in the indicated disk shelf voltage sensor. The shelf might be able to continue operation.

### Corrective Action

Verify that the power supply and the AC line are supplying power. Monitor the power grid for abnormalities. If all seems to be normal, replace the power supply. If the problem persists, contact NetApp technical support.

## Syslog Message

%s voltage warning for %s %d: %s%s. %s.

## Parameters

**prodChannel** (STRING): Disk shelf identity (channel, product, and disk shelf number).  
**typeText** (STRING): SCSI Enclosure Services (SES) type descriptor text for the power supply.  
**psNumber** (INT): SES element number of the power supply.  
**errorMsg** (STRING): Generic SES error.  
**errorText** (STRING): Description of the exact warning.  
**locationText** (STRING): Location of the module on the disk shelf.

## ses.system events

### ses.system.em.mmErr

#### Severity

EMERGENCY

#### Description

This message occurs when ONTAP® software does not support internal disk drives in this platform.

#### Corrective Action

Check with NetApp technical support whether the combination of this platform type with internal disks is supported on a later version of ONTAP. If so, upgrade to the appropriate ONTAP version.

## Syslog Message

Disk shelf %d containing a %.8s system with internal drives (product ID %.16s) is not a supported combination.

## Parameters

**ShelfID** (INT): ID of the disk shelf.  
**VendorID** (STRING): Vendor ID of the disk shelf.  
**ProductID** (STRING): Product ID of the disk shelf.

# ses.unsupported events

## ses.unsupported.shelf.psu

### Severity

EMERGENCY

### Description

This message occurs when a disk shelf has one unsupported power supply unit installed in it. The unsupported power fits correctly in the shelf, but its functional characteristics are incompatible with the supported power supply.

### Corrective Action

Replace all unsupported power supplies immediately. Use the "storage shelf show -psu-part-number" command to locate all unsupported power supplies on the disk shelf, using the part number provided in the message as the required command parameter. Contact NetApp technical support for assistance and part number for replacement power supplies.

## Syslog Message

Part number %s not supported on Product ID %s. %s: power supply %d. %s.

## Parameters

**partNumber** (STRING): Part number of the unsupported power supply.  
**ProdId** (STRING): Product identification of the shelf with the unsupported power supply.  
**prodChannel** (STRING): Disk shelf identity (channel and shelf number).  
**psNumber** (INT): SES element number of the failing power supply.  
**locationText** (STRING): Location of the unsupported power supply on the disk shelf.

## ses.unsupported.shelf.psus

### Severity

ERROR

### Description

This message occurs when a shelf has one or more unsupported power supplies installed in it. The unsupported power fits correctly in the shelf, but its functional characteristics are incompatible with the supported power supply.

### Corrective Action

Replace all unsupported power supplies immediately. Use the "storage shelf show -fields psu-part-number" command to and locate all unsupported power supplies on the disk shelf. Unsupported power supplies can be identified by the part number provided in the message. Contact NetApp technical support for assistance

and part number for replacement power supplies.

### **Syslog Message**

Part number %s not supported on Product ID %s. %s: %d unsupported power supplies.

### **Parameters**

**partNumber** (STRING): Part number of the unsupported power supply.

**ProdId** (STRING): Product identification of the shelf with unsupported power supply.

**prodChannel** (STRING): Disk shelf identity (channel and shelf number).

**psCount** (INT): Number of unsupported power supplies on the shelf.

## **ses.unsupported.shelves.psus**

### **Severity**

EMERGENCY

### **Description**

This message occurs when two or more shelves have unsupported power supply installed in it. The unsupported power fits correctly in the shelf, but its functional characteristics are incompatible with the supported power supply.

### **Corrective Action**

Replace all unsupported power supplies immediately. Use the "storage shelf show -fields psu-part-number" command to locate all unsupported power supplies on the disk shelf. Unsupported power supplies can be identified by the part number provided in the message. Contact NetApp technical support for assistance and part number for replacement power supplies.

### **Syslog Message**

Part number %s not supported on Product ID %s. %s and other shelves: %d unsupported power supplies.

### **Parameters**

**partNumber** (STRING): Part number of the unsupported power supply.

**ProdId** (STRING): Product identification of the shelf with unsupported power supply.

**prodChannel** (STRING): Disk shelf identity (channel and shelf number).

**psCount** (INT): Number of unsupported power supplies on the shelf.

# sfo events

## sfo.aggr events

### sfo.aggr.relocated.cu

#### Severity

INFORMATIONAL

#### Description

This message occurs when a storage failover aggregate is relocated from one node to another during a nondisruptive controller upgrade.

#### Corrective Action

(None).

#### Syslog Message

Aggregate %s was relocated from source node %s to destination node %s in %llu milliseconds during a nondisruptive controller upgrade. The source node has system ID %llu and the destination node has system ID %llu.

#### Parameters

**aggr\_name** (STRING): Name of the aggregate that was relocated.  
**src\_name** (STRING): Name of the source node.  
**dest\_name** (STRING): Name of the destination node.  
**time** (LONGINT): Time, in milliseconds, taken to relocate the aggregate.  
**src\_system\_id** (LONGINT): System ID of the source node.  
**dest\_system\_id** (LONGINT): System ID of the destination node.

### sfo.aggr.relocated.perm

#### Severity

INFORMATIONAL

#### Description

This message occurs when an SFO aggregate is "permanently" relocated from one node to another. However, even though the aggregate relocation is considered "permanent," you can still reverse or undo it.

#### Corrective Action

(None).

#### Syslog Message

Aggregate %s was permanently relocated from source node %s to destination node %s.

#### Parameters

**aggr\_name** (STRING): Name of the SFO aggregate that was relocated.  
**src\_name** (STRING): Name of the source node.  
**dest\_name** (STRING): Name of the destination node.  
**time** (LONGINT): Time, in milliseconds, taken to relocate the aggregate.

## sfo.aggr.relocated.temp

### Severity

INFORMATIONAL

### Description

This message occurs when a storage failover aggregate is temporarily relocated from one node to another. This usually happens when the aggregate is relocated during a performance optimized negotiated takeover.

### Corrective Action

(None).

### Syslog Message

Aggregate %s was temporarily relocated from source node %s to destination node %s in %llu milliseconds.

### Parameters

**aggr\_name** (STRING): Name of the aggregate that was relocated.

**src\_name** (STRING): Name of the source node.

**dest\_name** (STRING): Name of the destination node.

**time** (LONGINT): Time, in milliseconds, taken to relocate the aggregate.

## sfo.arl events

### sfo.arl.gb.failed.conversion

#### Severity

ERROR

#### Description

This message occurs when ARL/giveback of the Storage Failover (SFO) aggregate fails because of active volume conversion operations against one or more volumes on this aggregate while acquiring the RAID lock.

#### Corrective Action

When volume conversion is complete, re-run the command 'storage failover takeover/giveback'.

#### Syslog Message

%s %s.

#### Parameters

**msg** (STRING): Error that prevented ARL/GB.

**aggr** (STRING): Name of the aggregate that was not TO/GB.

### sfo.arl.notAllowed

#### Severity

NOTICE

#### Description

This message occurs when a requested Storage Failover (SFO) aggregate relocation operation is refused because the destination node is not the high-availability (HA) partner of the source node.

## Corrective Action

Verify that the destination node of the aggregate relocation is the source node's high-availability (HA) partner by using the "storage failover show" command. Retry the aggregate relocation with the correct source and destination node.

## Syslog Message

SFO aggregate relocation refused for aggregate %s because the destination is not an HA partner. Source node: %s (NVRAM ID: %llu), destination node: %s (NVRAM ID: %llu).

## Parameters

**aggr\_name** (STRING): Name of the aggregate.  
**src\_name** (STRING): Name of the source node.  
**src\_nvram\_id** (LONGINT): NVRAM ID of the source node.  
**dest\_name** (STRING): Name of the destination node.  
**dest\_nvram\_id** (LONGINT): NVRAM ID of the destination node.

## sfo.arl.recvdBadSignal

### Severity

INFORMATIONAL

### Description

This message occurs when the aggregate relocation worker thread receives an unexpected signal from the kernel subsystem.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**signal\_info** (STRING): Additional information on the signal received, indicating if it was a bad message type or a bad signal mask.  
**signal\_value** (INT): Value of the unexpected signal.

## sfo.giveback events

### sfo.giveback.aborted.bdfu

#### Severity

ERROR

#### Description

This message occurs when giveback of storage failover (SFO) aggregates is aborted because the background disk firmware update either on the source node or destination node took more than 120 seconds to finish.

#### Corrective Action

Rerun the "storage failover giveback" command.

## Syslog Message

Giveback is aborted: %s took more than 120 seconds to finish the background disk firmware update.

## Parameters

**nodename** (STRING): Node where the background firmware update took more than 120 seconds to finish.

## sfo.giveback.aggregatesAbort

### Severity

INFORMATIONAL

### Description

This message occurs when a giveback of non-root aggregates is aborted by either another takeover or other conflicting operation.

### Corrective Action

(None).

## Syslog Message

Giveback of non-root aggregates to node %s was aborted due to another takeover or other conflicting operation. "override-vetoes" was set to %s, "require-partner-waiting" was set to %s, and "only\_newly\_acquired" was set to %s.

## Parameters

**partnerName** (STRING): Partner name.

**override\_vetoes** (STRING): Flag that indicates whether the system overrides veto checks during a giveback operation. This flag corresponds to the "-override-vetoes" parameter of the "storage failover giveback" command. When set to true, it might result in some veto checks made by subsystems on the source node getting overridden.

**require\_partner\_waiting** (STRING): Flag that indicates whether, during a giveback, the storage is given back regardless of whether the partner node is available to take back the storage. This flag corresponds to the "-require-partner-waiting" parameter of the "storage failover giveback" command. When set to true, it might result in the giveback proceeding even if the destination node is not ready to receive the aggregate being migrated.

**only\_newly\_acquired** (STRING): Flag that indicates whether, during a giveback, only aggregates acquired during the last takeover will be given back.

## sfo.giveback.attemptExceeded

### Severity

ALERT

### Description

This message occurs when the number of attempts for automatic giveback of storage failover (SFO) aggregates exceeds the maximum allowed.

### Corrective Action

Use the "storage failover show-giveback" command to identify the subsystems that failed the giveback. Follow the corrective action specified in the EMS message displayed by those subsystems. Retry giveback by using the "storage failover giveback" command.

## Syslog Message

Attempts for automatic giveback of SFO aggregates exceeded the maximum number (%d) of allowed attempts.

## Parameters

**attempts** (INT): Maximum number of automatic giveback attempts allowed.

## sfo.giveback.complete

### Severity

INFORMATIONAL

### Description

This message occurs when a giveback of storage failover (SFO) aggregates to the SFO partner controller finishes.

### Corrective Action

(None).

## Syslog Message

Giveback of SFO aggregates to partner (%s) is complete. override-vetoes: %s, require-partner-waiting: %s, and only\_newly\_acquired: %s.

## Parameters

**partnerName** (STRING): Partner's name.

**override\_vetoes** (STRING): Flag that indicates whether the system overrides vetoes during a giveback operation. This flag corresponds to the "-override-vetoes" parameter of the "storage failover giveback" command. When set to true, it might result in some veto checks made by subsystems on the source node getting overridden.

**require\_partner\_waiting** (STRING): Flag that indicates whether, during a giveback, the storage is given back regardless of whether the partner node is available to take back the storage or not. This flag corresponds to the "-require-partner-waiting" parameter of the "storage failover giveback" command. When set to true, it might result in the giveback proceeding even if the destination node is not ready to receive the aggregate being migrated.

**only\_newly\_acquired** (STRING): Flag that indicates whether, during a giveback, only aggregates acquired during the last takeover will be given back.

## sfo.giveback.delayed.bdfu

### Severity

INFORMATIONAL

### Description

This message occurs when giveback of storage failover (SFO) aggregates is delayed because a background disk firmware update is in progress either on the source node or destination node. If the update does not finish within 120 seconds, giveback will be aborted.

### Corrective Action

(None).



## Syslog Message

(None).

## Parameters

**nodename** (STRING): Node where the background firmware update is in progress.

## sfo.giveback.failed

### Severity

ALERT

### Description

This message occurs when giveback of the partner's Storage Failover (SFO) aggregate fails.

### Corrective Action

Other error messages indicate the exact cause of failure along with the corrective action.

## Syslog Message

Giveback of aggregate %s failed due to %s.

## Parameters

**aggr** (STRING): Name of the aggregate that was not given back.

**msg** (STRING): Error that prevented giveback.

## sfo.giveback.OpFailed

### Severity

ALERT

### Description

This message occurs when the Storage Failover (SFO) phase of giveback fails.

### Corrective Action

Refer to other generated error messages to identify the exact cause of failure and the corrective action.

## Syslog Message

SFO phase of giveback failed by node %s because %s.

## Parameters

**node** (STRING): Name of the node that failed giveback.

**msg** (STRING): Error that prevented giveback.

## sfo.giveback.postMigrNotOk

### Severity

ALERT

### Description

This message occurs when the specified subsystem encounters an error during the post-migration phase of the giveback operation.

## Corrective Action

Check the event log for a subsystem-specific reason as to why the subsystem failed the post-migration phase of the giveback operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by that message.

## Syslog Message

The giveback operation of '%s' encountered an error during post-migration processing. The error was encountered by the subsystem '%s'.

## Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the giveback operation.  
**subsystem** (STRING): Name of the subsystem that failed the post-migration phase of the giveback operation.

## sfo.giveback.spares.failed

### Severity

ALERT

### Description

This message occurs when giveback of the partner's storage failover (SFO) spares fails.

### Corrective Action

Try the SFO giveback again.

### Syslog Message

Giveback of spares failed due to %s.

### Parameters

**msg** (STRING): Error that prevented giveback of SFO spares.

## sfo.ignoretakeover events

### sfo.ignoretakeover.envVarClearFailed

### Severity

ERROR

### Description

This message occurs when the environment indicating that mailbox takeover state should be rewritten could not be removed.

### Corrective Action

Boot into the firmware monitor and remove the ign-tkover environment variable.

### Syslog Message

Mailbox takeover state rewritten due to previous bringhome operation.

### Parameters

(None).

## **sfo.ignoretakeover.rewritingMailbox**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when mailbox takeover state is rewritten due to a previous bringhome operation.

### **Corrective Action**

None.

### **Syslog Message**

Mailbox takeover state rewritten due to previous bringhome operation.

### **Parameters**

(None).

## **sfo.partialgiveback events**

### **sfo.PartialGiveback**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a node is in partial giveback state because only some aggregates were given back during the giveback operation.

### **Corrective Action**

Use the "storage failover show-giveback" command to check the giveback status of the aggregates. This command displays the cause of giveback failure and suggests corrective actions. Then retry giveback by using the "storage failover giveback" command.

### **Syslog Message**

Node is in partial giveback state.

### **Parameters**

(None).

## **sfo.reassignfailed events**

### **sfo.reassignFailed**

### **Severity**

EMERGENCY

### **Description**

This message occurs during the Storage Failover takeover process, when the disks from the failed node cannot be reassigned due to the failure listed. Access to the failed node's volumes will be unavailable until this situation is fixed.

### **Corrective Action**

Use the "storage disk show -owner" command to verify that the surviving node can access the failed node's disks. Use the "storage disk show -owner" command to verify that the node can access the partner disks. Investigate and remedy the error listed. Use the "storage failover takeover" command to retry the takeover.

### **Syslog Message**

The takeover disk reassignment failed due to %s.

### **Parameters**

msg (STRING): Information about the error that prevented the reassignment.

## **sfo.reassignnpartner events**

### **sfo.reassignNoPartner**

#### **Severity**

ERROR

#### **Description**

This message occurs during the Storage Failover takeover process, when the disks from the failed node cannot be reassigned because they are inaccessible to the surviving node. Access to the failed node's volumes will be unavailable until this situation is fixed.

#### **Corrective Action**

Use the "storage disk show -owner" command to verify that the surviving node can access the failed node's disks. If the disks are still unavailable, check cabling and connectivity to the partner disks. If the disks are available, use the "storage failover takeover" command to reinitiate the takeover. Use the "storage disk show -owner" command to verify that the node can access the partner disks. When the node can access the partner disks, use the "storage failover takeover" command.

#### **Syslog Message**

The takeover disk reassignment failed because the partner disks were not accessible.

#### **Parameters**

(None).

## **sfo.retry events**

### **sfo.retry.autoGiveback**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when automatic giveback of one or more SFO aggregates fails. Giveback of these SFO aggregates will be retried after the specified time.

#### **Corrective Action**

(None).

## **Syslog Message**

Automatic giveback of SFO aggregates will be retried after %d minutes.

## **Parameters**

**minutes** (INT): Time until giveback is retried, in minutes.

# **sfo.retrying events**

## **sfo.retrying.autoGiveback**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when automatic giveback of SFO aggregates is retried.

### **Corrective Action**

(None).

### **Syslog Message**

Automatic giveback of SFO aggregates is being retried.

### **Parameters**

(None).

# **sfo.sendhome events**

## **sfo.sendhome.NoPartnerDisks**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the giveback request fails because there are no disks that belong to the partner.

### **Corrective Action**

Use the "storage disk show" command to validate that home information is set correctly on the disks that should belong to the partner.

### **Syslog Message**

Giveback request failed because there are no disks to return to the partner.

### **Parameters**

(None).

## **sfo.sendhome.partnerNoBoot**

### **Severity**

ERROR

## Description

This message occurs when the Storage Failover giveback process cannot proceed because the partner node did not boot successfully after its root aggregate was sent home. Volumes in the root aggregate are unavailable until you remedy this situation. Other aggregates belonging to the partner remain assigned to this node until you run another "storage failover giveback" command following a successful boot of the partner node.

## Corrective Action

Look at the partner logs and console output to determine why the node is not booting. Fix the problem, and then reboot the node. After the partner node boots successfully, run the "storage failover giveback" command again to return the remaining aggregates.

## Syslog Message

The giveback terminated because the partner node (%s) did not boot successfully.

## Parameters

**partner\_name** (STRING): Name of the partner node.

## sfo.sendhome.subsystemAbort

### Severity

ALERT

### Description

This message occurs when the specified subsystem aborts the Storage Failover giveback operation for the aggregate.

### Corrective Action

Check the syslog/EMS output for a subsystem-specific reason for aborting the giveback operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by the subsystem and then reissue the 'giveback' command. If you cannot perform the corrective action, then use the 'override-vetoes' option in the 'giveback' command to force the giveback.

### Syslog Message

The giveback operation of '%s' was aborted by '%s'.

### Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the giveback operation.

**subsystem** (STRING): Name of the subsystem that aborted the giveback operation.

## sfo.takenover events

### sfo.takenOver.aborted

### Severity

ERROR

### Description

This message occurs when an attempted optimized negotiated takeover is aborted by the node about to be taken over.

## Corrective Action

1. Use the "storage failover show" command to verify that the partner node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local and partner nodes are up. If they are not up, correct any network issues that could be preventing it.
3. Check whether any operation is running on the partner that might have slowed it down. If so, either abort the operation or wait until it is over.
4. (Optional) Use the "-bypass-optimization" parameter in the "storage failover takeover" command to bypass the optimized takeover by the partner node.
5. Use the "storage aggregate plex show" command to check mirror degraded aggregates, if the remote plex is offline, bring it online using the "storage aggregate plex online" command.

## Syslog Message

The local node has aborted an optimized negotiated takeover that was initiated by the partner. Reason: %s

## Parameters

**reason** (STRING): Reason for the failure, which might be as follows: 1. Conflicting takeover/giveback/relocation operation. 2. High-availability (HA) partner not being eligible for takeover. 3. Communication failure with the partner over the cluster network. 4. Local node has mirror degraded aggregates with remote plex inactive or unmirrored aggregates.

## sfo.takenOver.commFailed

### Severity

ALERT

### Description

This message occurs when the local node cannot communicate over the cluster network with its high-availability (HA) partner during an optimized negotiated takeover initiated by the HA partner node.

## Corrective Action

1. Use the "storage failover show" command to verify that the partner node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local and partner nodes are up. If they are not up, address any network issues that could be preventing it.
3. (Optional) Use the "-bypass-optimization" parameter in the "storage failover takeover" command to bypass the optimized takeover by the partner node.

## Syslog Message

The local node could not communicate with its HA partner, %s, (partner's D-blade ID: %s) during an optimized negotiated takeover initiated by the partner node. The communication request returned a status of %d. The level %d returned a status of %lld.

## Parameters

**partner\_name** (STRING): Name of the partner node.  
**partner\_dbladeid** (STRING): D-blade ID of the partner node.  
**status** (INT): Status returned by the hamsg message layer.  
**scope** (INT): Specific module level where the message did not succeed.  
**scope\_status** (LONGINT): Module-level specific reason the message did not succeed.

## sfo.takenOver.relocDone

### Severity

INFORMATIONAL

## Description

This message occurs during optimized negotiated takeover, when the high-availability (HA) partner notifies the local node that it has completed the relocation of all aggregates to the local node.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**partner\_name** (STRING): Name of the partner node.

**partner\_sysid** (LONGINT): System ID of the partner node.

**relocation\_status** (INT): Status of the "relocation done" response received from the HA partner.

# sfo.takeover events

## sfo.takeover.aborted

### Severity

ERROR

### Description

This message occurs when the taking over node cannot complete an attempted optimized negotiated takeover. Any data LIFs that were migrated as part of the failed takeover operation will not be auto-reverted back.

### Corrective Action

1. Use the "storage failover show" command to verify that the partner node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local and partner nodes are up. If they are not up, correct any network issues that could be preventing it.
3. Check whether any operation is running on the partner that might have slowed it down. If so, either abort the operation or wait until it is over.
4. (Optional) Use the "-bypass-optimization" parameter in the "storage failover takeover" command to bypass the optimized takeover by the local node.

### Syslog Message

An optimized negotiated takeover of the HA partner node by the local node has been aborted. Reason: %s

### Parameters

**reason** (STRING): Reason for the failure, which might be as follows: 1. Conflicting takeover/giveback/relocation operation. 2. High-availability (HA) partner not being eligible for takeover. 3. Communication failure with the partner over the cluster network.

## sfo.takeover.aggrProcessTime

### Severity

NOTICE

### Description

This message occurs during the planned negotiated takeover of an aggregate, to indicate the time taken by a particular process during that takeover.



## Corrective Action

(None).

## Syslog Message

Time taken for %s during takeover of the aggregate '%s' was %llu milliseconds.

## Parameters

**process** (STRING): Description of the process.

**aggr** (STRING): Name of the aggregate being taken over.

**time** (LONGINT): Time taken by the specified process in milliseconds.

## sfo.takeover.aggrSubsysTime

### Severity

NOTICE

### Description

This message occurs during the planned negotiated takeover of an aggregate, to indicate the time taken by a subsystem to perform a particular process during that takeover.

### Corrective Action

(None).

### Syslog Message

Subsystem %s took %llu milliseconds to complete %s during planned negotiated takeover of the aggregate %s.

### Parameters

**subsystem** (STRING): Name of the subsystem performing the takeover operation.

**time** (LONGINT): Time taken by the subsystem, in milliseconds, to complete the specified process.

**process** (STRING): Description of the process.

**aggr** (STRING): Name of the aggregate being taken over.

## sfo.takeover.bypassed

### Severity

INFORMATIONAL

### Description

This message occurs when a negotiated takeover bypasses the negotiated takeover optimization.

### Corrective Action

(None).

### Syslog Message

Negotiated takeover of node %s with system-ID %u has bypassed the negotiated takeover optimization.  
Reason: %s

### Parameters

**partner\_name** (STRING): Name of the partner node.

**partner\_sysid** (LONGINT): System ID of the partner node.

**reason** (STRING): Reason that negotiated takeover optimization was bypassed.

## **sfo.takeover.commExpired**

### **Severity**

ERROR

### **Description**

This message occurs during optimized negotiated takeover, when the local node does not receive a response within a designated time period, while communicating over the cluster network with its high-availability (HA) partner.

### **Corrective Action**

1. Check whether any operation is running on the partner node that might have slowed it down, such as another aggregate relocation. Also check for conflicting operations, such as another takeover, giveback, or shutdown. If so, either abort the operation or wait until it is over.
2. (Optional) Use the "-bypass-optimization" parameter in the "storage failover takeover" command to bypass the optimized takeover by the local node.

### **Syslog Message**

The local node timed out while waiting for a response from its HA partner, %s, over the cluster network.

### **Parameters**

**partner\_name** (STRING): Name of the partner node.

**partner\_dbladeid** (STRING): D-blade ID of the partner node.

## **sfo.takeover.commFailed**

### **Severity**

ALERT

### **Description**

This message occurs when the local node cannot communicate over the cluster network with its high-availability (HA) partner during an optimized negotiated takeover initiated by the local node.

### **Corrective Action**

1. Use the "storage failover show" command to verify that the partner node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local and partner nodes are up. If they are not up, address any network issues that could be preventing it.
3. (Optional) Use the "-bypass-optimization" parameter in the "storage failover takeover" command to bypass the optimized takeover by the local node.

### **Syslog Message**

The local node could not communicate with its HA partner, %s, (partner's D-blade ID: %s) during an optimized negotiated takeover initiated by the local node. The communication request returned a status of %d. The level %d returned a status of %lld.

### **Parameters**

**partner\_name** (STRING): Name of the partner node.

**partner\_dbladeid** (STRING): D-blade ID of the partner node.

**status** (INT): Status returned by the hamsg message layer.

**scope** (INT): Specific module level where the message did not succeed.

**scope\_status** (LONGINT): Module-level specific reason the message did not succeed.

## **sfo.takeover.noAggrReloc**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when no SFO aggregate was relocated by the high-availability (HA) partner during optimized negotiated takeover.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**partner\_name** (STRING): Name of the partner node.

**partner\_sysid** (LONGINT): System ID of the partner node.

## **sfo.takeover.relocDone**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs during optimized negotiated takeover, when the local node receives a notification from its high-availability (HA) partner node stating that the partner has completed the relocation of all aggregates to the local node.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**partner\_name** (STRING): Name of the partner node.

**partner\_sysid** (LONGINT): System ID of the partner node.

**relocation\_status** (INT): Status of the "relocation done" response received from the HA partner.

## **sfo.takeover.sfoStart**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the node doing the takeover starts the SFO phase of the optimized negotiated takeover.

### Corrective Action

(None).

### Syslog Message

The local node has initiated optimized negotiated takeover of its partner (%s).

### Parameters

**partner\_name** (STRING): Name of the partner node.

**partner\_sysid** (LONGINT): System ID of the partner node.

## sfo.tkabort events

### sfo.tkAbort.ca.check.failed

#### Severity

ERROR

#### Description

This message occurs when the destination node cannot contact the object stores during the planned takeover of a FabricPool aggregate.

#### Corrective Action

Verify that the intercluster LIF is online and functional by using the "network interface show" command. Also, check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. Additionally, verify that the configuration of the object store has not changed, and that login and connectivity information is still accurate by using the "aggregate object-store config show" command.

#### Syslog Message

Planned SFO takeover of aggregate '%s' (uuid: %s) failed due to %s.

#### Parameters

**vol** (STRING): Name of the aggregate.

**aggr\_uuid** (STRING): UUID of the aggregate.

**reason** (STRING): Reason for the prevention of the planned takeover operation.

## sfo.tkovertoclusterwait events

### sfo.tkovertoclusterwait

#### Severity

NOTICE

#### Description

This message occurs when a node whose M-host applications have not fully stabilized initiates a takeover followed by a giveback to its high-availability (HA) partner. After the giveback, this node's Failover Monitor state indicates that M-host stabilization is not yet complete.

#### Corrective Action

Contact NetApp technical support to diagnose the issue further.

**Syslog Message**

(None).

**Parameters**

(None).

# sfu events

## sfu.adaptersuspendio events

### sfu.adapterSuspendIO

#### Severity

INFORMATIONAL

#### Description

This message occurs during a disk shelf firmware update on a disk shelf that cannot perform I/O while updating firmware. Typically, the shelves involved are bridge-based.

#### Corrective Action

I/O resumes automatically after the firmware update is complete.

#### Syslog Message

Suspending I/O to %s adapter %s for %d seconds while shelf firmware is updated.

#### Parameters

**adapterType** (STRING): Type of adapter that is updating firmware.

**channelName** (STRING): Name of the adapter on which shelves will be updating firmware.

**minutes** (INT): Amount of time that I/O will be suspended on the adapter.

### sfu.adapterSuspendIO.ndu

#### Severity

INFORMATIONAL

#### Description

This message occurs during a disk shelf firmware update on a disk shelf that can perform I/O, but might have to suspend management requests while updating firmware. Typically, the shelves involved are DS12-ESAS, DS20-ESAS, DS4243.

#### Corrective Action

Suspended data traffic resumes automatically after the firmware update is complete.

#### Syslog Message

Suspending %s to %s adapter %s for %d seconds while shelf firmware is updated.

#### Parameters

**dataTrafficType** (STRING): Type of data traffic: SMP.

**adapterType** (STRING): Type of adapter that is updating firmware.

**channelName** (STRING): Name of the adapter on which shelves will be updating firmware.

**minutes** (INT): Amount of time that specified data traffic will be suspended on the adapter.

## sfu.auto events

## sfu.auto.update.off.impact

### Severity

ERROR

### Description

This message occurs when ONTAP® software cannot automatically upgrade the firmware on a disk shelf because the hidden 'shelf.fw.auto.update' option is set to off.

### Corrective Action

Set the hidden option 'shelf.fw.auto.update' to "on", or update the firmware manually.

### Syslog Message

Firmware could not be updated on I/O module %s with firmware version %s on %s disk shelf %s.shelf%d.  
Update the firmware manually.

### Parameters

**moduleType** (STRING): I/O controller module type, for example, AT-FCX/ESH4/SAS/IOM3.

**productRev** (STRING): Firmware revision level.

**shelfType** (STRING): Disk shelf enclosure type, for example, DS14-Mk4-FC, DS14-Mk2-FC, or DS14-Mk2-AT.

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf enclosure.

**shelfId** (INT): ID of the disk shelf enclosure.

## sfu.badbridgeshelfconfig events

### sfu.badBridgeShelfConfig

### Severity

ALERT

### Description

This message occurs when the disk shelf firmware update aborts because it cannot contact the two peer FC-to-SAS bridges.

### Corrective Action

Make sure that the storage controller is correctly cabled to the two peer FC-to-SAS bridges and that the bridges are correctly cabled to the SAS disk shelves. Then retry the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

Firmware update aborted on disk shelf %s.shelf%d module %s. Unable to contact two peer FC-to-SAS bridges.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf that aborted the firmware update.

**shelfId** (INT): ID of the disk shelf that aborted the firmware update.

**module** (STRING): Module that aborted the firmware update.

# sfu.controllerElementsPerShelf events

## sfu.controllerElementsPerShelf

### Severity

INFORMATIONAL

### Description

This message occurs when a disk shelf firmware download determines the number of controller elements per shelf that can be downloaded.

### Corrective Action

(None).

### Syslog Message

%s: %d ES controller %s can be updated on %s.shelf%d.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**count** (INT): Number of ES controller elements on this shelf that can be downloaded.

**pluralElement** (STRING): Element or Elements.

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware can be downloaded.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware can be downloaded.

# sfu.download events

## sfu.download.details

### Deprecated

Deprecated as of ONTAP® version 9.6. A new EMS message is added.

### Severity

NOTICE

### Description

This message occurs when a disk shelf firmware update completes download firmware to a disk shelf or shelves in the system.

### Corrective Action

Shelf firmware update will be attempted again in the next background shelf firmware update.

### Syslog Message

%s: Download of firmware file %s on %s.shelf%d might not have completed on one module as %s. An attempt to download again will be done in the next cycle.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware download error occurred.



**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download error occurred.  
**ErrorReason** (STRING): More information about the reason for download error.

## sfu.download.done.details

### Severity

NOTICE

### Description

This message occurs when a shelf firmware update finishes downloading firmware to a shelf or shelves in the system.

### Corrective Action

Shelf firmware update will be attempted again in the next background shelf firmware update.

### Syslog Message

%s: Download of firmware file on %s.shelf%d might not have been completed on one module because %s.  
An attempt to download again will be done in the next cycle.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".  
**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware download error occurred.  
**shelfId** (INT): ID of the shelf on which the shelf firmware download error occurred.  
**ErrorReason** (STRING): More information about the reason for download error.

## sfu.downloadctrllerbridge events

### sfu.downloadCtrllerBridge

### Severity

INFORMATIONAL

### Description

This message occurs when a disk shelf firmware download starts on a particular disk shelf.

### Corrective Action

(None).

### Syslog Message

%s: Downloading %s on ES controller module %s via SES device %s.%d on %s.shelf%d.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".  
**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.  
**module** (STRING): Either Module A or Module B, in which Module A is the upper module in the enclosure on the back side and Module B is the lower module in the enclosure on the back side.  
**channelNameSES** (STRING): Name of the adapter port or switch port on the SES services device on which the disk shelf firmware download is occurring.  
**scsild** (INT): SCSI ID of the SES services device used to download disk shelf firmware.  
**channelNameShelf** (STRING): Name of the adapter port or switch port on which the disk shelf firmware

download is occurring.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download is occurring.

## sfu.downloadererror events

### sfu.downloadError

#### Severity

ERROR

#### Description

This message occurs when a disk shelf firmware update fails to successfully download firmware to a disk shelf or shelves in the system.

#### Corrective Action

Redownload the latest disk shelf firmware from the NetApp Support Site at <http://support.netapp.com/NOW/download/tools/diskshelf/> and attempt to download disk shelf firmware again, using the (advanced privilege) "storage shelf firmware update" command.

#### Syslog Message

%s: Error downloading firmware file %s on %s.shelf%d.%s

#### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware download error occurred.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download. error occurred.

**ErrorReason** (STRING): More information about the reason for download error.

## sfu.downloadingcontroller events

### sfu.downloadingController

#### Severity

INFORMATIONAL

#### Description

This message occurs when a disk shelf firmware download starts on a specific disk shelf.

#### Corrective Action

(None).

#### Syslog Message

%s: Downloading %s on disk shelf controller module %s on %s.shelf%d.

#### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**module** (STRING): Module on which the disk shelf firmware download is occurring.

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware

download is occurring.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download is occurring.

## sfu.downloadingctrlerr1xx events

### sfu.downloadingCtrllerR1XX

#### Severity

INFORMATIONAL

#### Description

This message occurs when a disk shelf firmware download starts on a particular disk shelf.

#### Corrective Action

(None).

#### Syslog Message

%s: Downloading %s on ES controller on %s.shelf%d.

#### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware download is occurring.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download is occurring.

## sfu.downloadstarted events

### sfu.downloadStarted

#### Severity

INFORMATIONAL

#### Description

This message occurs when a disk shelf firmware update starts to download disk shelf firmware.

#### Corrective Action

(None).

#### Syslog Message

Update of disk shelf firmware started on %d %s.

#### Parameters

**count** (INT): Number of shelves that are starting disk shelf firmware download.

**pluralShelf** (STRING): Shelf or Shelves.

## sfu.downloadsuccess events

## sfu.downloadSuccess

### Severity

INFORMATIONAL

### Description

This message occurs when disk shelf firmware updates successfully.

### Corrective Action

(None).

### Syslog Message

%s: Firmware file %s downloaded on %s.shelf%d.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**channelName** (STRING): Name of the adapter port or switch port on which the shelf firmware download occurred.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download occurred.

## sfu.downloadsummary events

### sfu.downloadSummary

#### Severity

INFORMATIONAL

#### Description

This message occurs when a disk shelf firmware update completes successfully.

#### Corrective Action

(None).

#### Syslog Message

Shelf firmware updated on %d %s.

#### Parameters

**shelfCount** (INT): Number of shelves that successfully updated.

**pluralShelf** (STRING): Shelf or Shelves.

## sfu.downloadsummaryerrors events

### sfu.downloadSummaryErrors

#### Severity

ERROR

## Description

This message occurs when a disk shelf firmware update completes without successfully downloading to all shelves it attempted.

## Corrective Action

Issue the (advanced privilege) "storage shelf firmware update" command again.

## Syslog Message

Shelf firmware updated on %d %s, failed to update on %d %s.

## Parameters

**successCount** (INT): Number of shelves successfully updated.

**pluralShelf** (STRING): Shelf or Shelves.

**failedCount** (INT): Number of shelves that failed to update.

**pluralShelfFailed** (STRING): Shelf or Shelves.

# sfu.fcdownloadfailed events

## sfu.FCDownloadFailed

### Severity

ERROR

### Description

This message occurs when a disk shelf firmware update fails to download shelf firmware to an FC or an ATA shelf successfully.

### Corrective Action

Redownload the latest disk shelf firmware from the NetApp Support Site at <http://support.netapp.com/NOW/download/tools/diskshelf/> and attempt to download disk shelf firmware again by using the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

%s: Download failed on %s on ES controller module %s via SES device %s.%d on %s.shelf%d.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of file in /etc/shelf\_fw that was being downloaded.

**module** (STRING): Module A is the upper module in the enclosure on the back side. Module B is the lower module in the enclosure on the back side.

**channelNameSES** (STRING): Name of the adapter port or switch port on the SES services device on which the disk shelf firmware download failed.

**scsild** (INT): SCSI ID of the SES services device used to download disk shelf firmware.

**channelNameShelf** (STRING): Name of the adapter port or switch port on which the disk shelf firmware download failed.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download failed.

# sfu.firmwaredownrev events

## sfu.firmwareDownrev

### Severity

ERROR

### Description

This message occurs when disk shelf firmware is downrev and therefore cannot be updated automatically. Before updating the disk shelf firmware, always check the NetApp Support Site for the latest disk shelf firmware available: <http://support.netapp.com/NOW/download/tools/diskshelf/>.

### Corrective Action

Copy updated disk shelf firmware into the /etc/shelf\_fw directory on the storage appliance and then manually issue the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

Disk shelf firmware needs to be updated on %d disk %s.

### Parameters

**downrevShelves** (INT): Number of shelves in the system that are downrev.  
**pluralShelf** (STRING): Shelf or Shelves.

## sfu.firmwareDownrev.shelf

### Severity

ERROR

### Description

This message occurs when the disk shelf is running downrev firmware, and a later version is available in the /etc/shelf\_fw directory. The shelf cannot be updated automatically, possibly because automatic shelf firmware update is disabled.

### Corrective Action

Use the "options shelf.fw.auto.update" command to check whether automatic shelf firmware update is off. Manually issue the (advanced privilege) "storage shelf firmware update" command if background update is disabled.

### Syslog Message

Shelf %s has downrev firmware.

### Parameters

**ShelfId** (STRING): Shelf that is running downrev firmware.

## sfu.firmwareuptodate events

### sfu.firmwareUpToDate

### Severity

INFORMATIONAL

## Description

This message occurs when a disk shelf firmware update is requested but all shelves are determined to be updated already to the latest version of firmware available on the system. Before updating the disk shelf firmware, always check the NetApp Support Site for the latest disk shelf firmware available: <http://support.netapp.com/NOW/download/tools/diskshelf/>.

## Corrective Action

(None).

## Syslog Message

Firmware is up-to-date on all disk shelves.

## Parameters

(None).

# sfu.nsmdownloadfailed events

## sfu.NSMDownloadFailed

### Severity

ERROR

### Description

This message occurs when a disk shelf firmware update fails to download disk shelf firmware to a disk shelf successfully.

### Corrective Action

Go to the NetApp Support Site at <http://support.netapp.com/NOW/download/tools/diskshelf/> and try downloading the disk shelf firmware again by using the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

%s: Download failed on %s on disk shelf controller module %s on disk shelf %s.%d.

### Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**module** (STRING): Module on which the disk shelf firmware download failed.

**channelNameSES** (STRING): Name of the adapter port or switch port on the SCSI enclosure services (SES) device on which the disk shelf firmware download failed.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download failed.

# sfu.partnerinaccessible events

## sfu.partnerInaccessible

### Severity

ALERT

## Description

This message occurs in a high-availability (HA) configuration in which communication between partner nodes cannot be established.

## Corrective Action

Verify that the HA interconnect is operational. Retry the (advanced privilege) "storage shelf firmware update" command.

## Syslog Message

HA partner is inaccessible. Shelf firmware update canceled.

## Parameters

(None).

# sfu.partnernotresponding events

## sfu.partnerNotResponding

### Severity

ALERT

### Description

This message occurs when one node in a high-availability (HA) configuration does not respond to firmware download requests from another node, in which case the other node cannot download disk shelf firmware.

### Corrective Action

Make certain that the HA interconnect is up and running on both nodes of the HA pair and then attempt to re-download disk shelf firmware, using the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

Partner either responded in the negative, or did not respond in %d seconds. Aborting shelf firmware update.

### Parameters

**timeout** (INT): Number of seconds the timeout is set for.

# sfu.partnerrefusedupdate events

## sfu.partnerRefusedUpdate

### Severity

ALERT

### Description

This message occurs in a high-availability (HA) configuration in which one node refuses firmware download requests from its partner node, in which case the partner node cannot download disk shelf firmware.

### Corrective Action

Make certain that the HA partners are both running the same version of ONTAP® software and that the HA interconnect is up and running. Attempt the (advanced privilege) "storage shelf firmware update" command again.



## Syslog Message

Partner node refused permission to update shelf firmware.

## Parameters

(None).

# sfu.partnerupdatecomplete events

## sfu.partnerUpdateComplete

### Severity

INFORMATIONAL

### Description

This message occurs in a high-availability (HA) configuration in which an HA partner was downloading disk shelf firmware and the download was completed, at which point this notification is sent and enclosure services are resumed by the partner.

### Corrective Action

(None).

## Syslog Message

Partner is no longer updating disk shelf firmware. Resuming enclosure services.

## Parameters

(None).

# sfu.partnerupdatetimeout events

## sfu.partnerUpdateTimeout

### Severity

INFORMATIONAL

### Description

This message occurs in a high-availability (HA) configuration in which an HA partner was downloading disk shelf firmware and the download has timed out. The system resumes enclosure services monitoring.

### Corrective Action

Verify that the HA interconnect is operational and then retry the (advanced privilege) "storage shelf firmware update" command.

## Syslog Message

Waited %d minutes for partner shelf firmware update. Resuming enclosure services.

## Parameters

**minutes** (INT): Duration of the partner update timeout, in minutes.

# sfu.rebootrequest events

## sfu.rebootRequest

### Severity

INFORMATIONAL

### Description

This message occurs when the disk shelf firmware update is complete and the disk shelf will be rebooted to run the new code.

### Corrective Action

(None).

### Syslog Message

Issuing a request to reboot disk shelf %s.shelf%d module %s.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf to be rebooted.

**shelfId** (INT): ID of the disk shelf to be rebooted.

**module** (STRING): Module to be rebooted.

# sfu.rebootrequestfailure events

## sfu.rebootRequestFailure

### Severity

ERROR

### Description

This message occurs when an attempt to issue a reboot request after downloading disk shelf firmware fails, indicating a software error.

### Corrective Action

Retry the (advanced privilege) "storage shelf firmware update" command.

### Syslog Message

Could not issue the request to reboot disk shelf %s.shelf%d module %s.

### Parameters

**channelName** (STRING): Name of the adapter port or switch port of the disk shelf to be rebooted.

**shelfId** (INT): ID of the disk shelf to be rebooted.

**module** (STRING): Module to be rebooted.

# sfu.resumediskio events

## sfu.resumeDiskIO

**Severity**

INFORMATIONAL

**Description**

This message occurs when a disk shelf firmware update is completed and disk I/O is resumed.

**Corrective Action**

(None).

**Syslog Message**

Resuming disk I/O after shelf firmware update.

**Parameters**

(None).

## sfu.sasdownloadfailed events

### sfu.SASDownloadFailed

**Severity**

ERROR

**Description**

This message occurs when a disk shelf firmware update fails to download disk shelf firmware to a disk shelf successfully.

**Corrective Action**

Go to the NetApp Support Site at <http://support.netapp.com/NOW/download/tools/diskshelf/> and try downloading the disk shelf firmware again by using the (advanced privilege) "storage shelf firmware update" command.

**Syslog Message**

%s: Download failed on %s on disk shelf controller module %s on disk shelf %s.%d.

**Parameters**

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**fileName** (STRING): Name of the file in /etc/shelf\_fw that was being downloaded.

**module** (STRING): Module on which the disk shelf firmware download failed.

**channelNameSES** (STRING): Name of the adapter port or switch port on the SCSI enclosure services (SES) device on which the disk shelf firmware download failed.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware download failed.

## sfu.statuscheckfailure events

### sfu.statusCheckFailure

**Severity**

ERROR

## Description

This message occurs when the (advanced privilege) "storage shelf firmware update" command encounters a failure while attempting to read the status of the firmware update in progress.

## Corrective Action

Retry the (advanced privilege) "storage shelf firmware update" command.

## Syslog Message

%s: failed status check of f/w update on %s.shelf%d [%x] %s, retry\_count = %d, %s.

## Parameters

**cmdName** (STRING): Name of the command issued, such as "storage shelf firmware update".

**channelName** (STRING): Name of the adapter port or switch port on which the disk shelf firmware is to be downloaded.

**shelfId** (INT): ID of the disk shelf on which the disk shelf firmware is to be downloaded.

**outStatusFlags** (INTHEX): Output status flags value of the I/O operation (ONTAP® software internal value).

**ModuleType** (STRING): Module type, such as AT-FCX, AT-FC, ESH, LRC, and so on.

**retryCount** (INT): Number of times the status read has been retried.

**DeviceStatus** (STRING): Whether the SES target device status is OK or MISSING.

# sfu.suspenddiskio events

## sfu.suspendDiskIO

### Severity

INFORMATIONAL

### Description

This message occurs when a disk shelf firmware update is started and disk I/O is suspended.

### Corrective Action

(None).

### Syslog Message

Suspending disk I/O to reboot shelf modules for shelf firmware update.

### Parameters

(None).

# sfu.suspendses events

## sfu.suspendSES

### Severity

INFORMATIONAL

### Description

This message occurs when a disk shelf firmware update is requested in a high-availability (HA) environment. One partner node updates the firmware on the disk shelf module while the HA partner node

temporarily disables enclosure services while the firmware update is in process.

**Corrective Action**

(None).

**Syslog Message**

Suspending enclosure services—partner is updating disk shelf firmware.

**Parameters**

(None).

# shelf events

## shelf.config events

### shelf.config.mixed

**Severity**

NOTICE

**Description**

This message occurs at system boot time, when a stand-alone system's storage configuration is a mixture of redundant paths and single paths.

**Corrective Action**

(None).

**Syslog Message**

System is using a mixture of single path, multi-path or quad-path storage configurations.

**Parameters**

(None).

### shelf.config.mixedha

**Severity**

NOTICE

**Description**

This message occurs at system boot time, when an HA system's storage configuration is a mixture of redundant paths and single paths.

**Corrective Action**

(None).

**Syslog Message**

System is using a mixture of single path, multi-path or quad-path HA storage configurations.

**Parameters**

(None).

### shelf.config.mpha

**Severity**

INFORMATIONAL

**Description**

This event is generated at the Filer boot time, when it is a HA configuration and all the storage shelf enclosures accessible by the Filer have redundant paths. This message is for informational purpose only.

**Corrective Action**

(None).

**Syslog Message**

All attached storage on the system is multi-pathed HA.

**Parameters**

(None).

**shelf.config.multipath****Severity**

INFORMATIONAL

**Description**

This event is generated at the Filer boot time, when it is a stand alone configuration and all the storage shelf enclosures accessible by the Filer have redundant paths. This message is for informational purpose only.

**Corrective Action**

(None).

**Syslog Message**

All attached storage on the system is multi-pathed.

**Parameters**

(None).

**shelf.config.qpha****Severity**

NOTICE

**Description**

This message occurs at system boot time, when an HA system's storage configuration is quad-path.

**Corrective Action**

(None).

**Syslog Message**

System storage is in a quad-path HA configuration.

**Parameters**

(None).

**shelf.config.quadpath****Severity**

NOTICE

**Description**

This message occurs at system boot time, when a stand-alone system's storage configuration is quad-path.

**Corrective Action**

(None).

**Syslog Message**

System storage is in a quad-path configuration.

**Parameters**

(None).

**shelf.config.single****Severity**

INFORMATIONAL

**Description**

This event is generated at the Filer boot time, when it is a stand alone configuration and all the storage shelf enclosures accessible by the Filer have only one active path. This message is for informational purpose only.

**Corrective Action**

(None).

**Syslog Message**

System is using single path attached storage only.

**Parameters**

(None).

**shelf.config.spha****Severity**

INFORMATIONAL

**Description**

This event is generated at the Filer boot time, when it is a HA configuration and all the storage shelf enclosures accessible by the Filer have only one active path. This message is for informational purpose only.

**Corrective Action**

(None).

**Syslog Message**

System is using single path HA attached storage only.

**Parameters**

(None).



## shelf.config.tomixed

### Severity

NOTICE

### Description

This message occurs in a stand-alone system when the storage configuration changes from a single-path, multi-path, or quad-path configuration, to a mixed path configuration.

### Corrective Action

(None).

### Syslog Message

System has transitioned to a mixture of single path, multi-path or quad-path storage configurations.

### Parameters

(None).

## shelf.config.tomixedha

### Severity

NOTICE

### Description

This message occurs in an HA system when the storage configuration changes from a single-path HA, multi-path HA, or quad-path HA configuration to a mixed path HA configuration.

### Corrective Action

(None).

### Syslog Message

System has transitioned to a mixture of single, multi-path or quad-path storage configurations.

### Parameters

(None).

## shelf.config.tompha

### Severity

INFORMATIONAL

### Description

This event is generated by the periodic storage health monitor process in a HA Filer configuration, when the Filer storage configuration changes from the single path HA configuration to the multi-path HA configuration.

### Corrective Action

(None).

### Syslog Message

System has transitioned to multi-path HA attached storage

**Parameters**

(None).

**shelf.config.tomultipath****Severity**

INFORMATIONAL

**Description**

This event is generated by the periodic storage health monitor process in a stand alone Filer configuration, when the Filer storage configuration changes from the single-path configuration to the multi path configuration.

**Corrective Action**

(None).

**Syslog Message**

System has transitioned to multi-path attached storage

**Parameters**

(None).

**shelf.config.toqpha****Severity**

NOTICE

**Description**

This message occurs in an HA system when the storage configuration changes to quad-path.

**Corrective Action**

(None).

**Syslog Message**

System has transitioned to a quad-path storage configuration.

**Parameters**

(None).

**shelf.config.toquadpath****Severity**

NOTICE

**Description**

This message occurs in a stand-alone system. System has transitioned to a quad-path storage configuration.

**Corrective Action**

(None).

**Syslog Message**

System has transitioned to quad-path attached storage

**Parameters**

(None).

**shelf.config.tosingle****Severity**

INFORMATIONAL

**Description**

This event is generated by the periodic storage health monitor process in a stand alone Filer configuration, when the Filer storage configuration changes from the multi-path configuration to the single path configuration.

**Corrective Action**

(None).

**Syslog Message**

System has transitioned to single path attached storage

**Parameters**

(None).

**shelf.config.tospha****Severity**

INFORMATIONAL

**Description**

This event is generated by the periodic storage health monitor process in a HA Filer configuration, when the Filer storage configuration changes from the multi-path HA configuration to the single path HA configuration.

**Corrective Action**

(None).

**Syslog Message**

System has transitioned to single path HA attached storage

**Parameters**

(None).

**shelf.config.toTripath****Severity**

NOTICE

**Description**

This message occurs when a stand-alone system changes to a tri-path storage configuration.

**Corrective Action**

(None).

**Syslog Message**

System transitions to a tri-path storage configuration.

**Parameters**

(None).

**shelf.config.toTripathHA****Severity**

NOTICE

**Description**

This message occurs in a high-availability (HA) system when the storage configuration changes to tri-path.

**Corrective Action**

(None).

**Syslog Message**

System transitions to a tri-path HA storage configuration.

**Parameters**

(None).

**shelf.config.tripath****Severity**

NOTICE

**Description**

This message occurs at system boot time, when a stand-alone system's storage configuration is tri-path.

**Corrective Action**

(None).

**Syslog Message**

System storage is in a tri-path configuration.

**Parameters**

(None).

**shelf.config.tripathHA****Severity**

NOTICE

## Description

This message occurs at system boot time, when a high-availability (HA) system's storage configuration is tri-path.

## Corrective Action

(None).

## Syslog Message

System storage is in a tri-path HA storage configuration.

## Parameters

(None).

# shelf.identify events

## shelf.identify.start

### Severity

INFORMATIONAL

### Description

This message occurs when the user turns on the disk shelf's Identify mode.

### Corrective Action

(None).

### Syslog Message

Disk shelf %s has entered Identify mode (Shelf Attention Code: %s)

### Parameters

**shelfName** (STRING): Disk shelf name.

**panelCode** (STRING): Code displayed on the disk shelf's 7-segment LED OPS panel.

## shelf.identify.stop

### Severity

INFORMATIONAL

### Description

This message occurs when the user turns off a disk shelf's Identify mode.

### Corrective Action

(None).

### Syslog Message

Disk shelf %s has exited Identify mode.

### Parameters

**shelfName** (STRING): Disk shelf name.

# shelf.io events

## shelf.IO.in.progress

### Severity

NOTICE

### Description

This message occurs when a request cannot be performed because I/O is still in progress.

### Corrective Action

(None).

### Syslog Message

Could not perform the request; I/O to disk shelf is still in progress.

### Parameters

(None).

# shelf.module events

## shelf.module.configuration.failure

### Severity

ERROR

### Description

This message occurs when the system cannot complete the configuration of a disk shelf module with the desired parameters.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Could not complete the configuration of disk shelf module %s on channels %s, disk shelf ID %d.

### Parameters

**element** (STRING): Disk shelf module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID of the disk shelf.

## shelf.module.element.missing

### Severity

ALERT

### Description

This message occurs when an element is missing, unused, or disconnected.

### Corrective Action

Verify that both disk shelf modules are fully seated and locked in place. Also make sure that they are of the same product family.

### Syslog Message

Disk shelf module %s is MISSING on channels %s disk shelf ID %d.

### Parameters

**element** (STRING): Disk shelf module designator (A is the upper module, B is the lower module).

**channels** (STRING): Channels that have the error if local, PARTNER for high-availability partner, UNUSED if unconnected.

**shelfID** (INT): Disk shelf ID of the disk shelf reporting the error.

## shelf.reset events

### shelf.reset.counters

#### Severity

INFORMATIONAL

#### Description

This message occurs when ONTAP® software resets the error counters on a disk shelf module.

#### Corrective Action

(None).

#### Syslog Message

Resetting error counters for module %s in disk shelf %s.%d.

#### Parameters

**element** (STRING): Disk shelf module designator (A is the upper module, B is the lower module).

**channel** (STRING): Channel the disk shelf module is on.

**shelfID** (INT): ID of the disk shelf the module is in.

## shelf.test events

### shelf.test.in.progress

#### Severity

INFORMATIONAL

#### Description

This message occurs when a TEST request is now in progress.

#### Corrective Action

(None).

#### Syslog Message

Test request in progress.

**Parameters**

(None).

**shelf.test.request.error****Severity**

ERROR

**Description**

This message occurs when a TEST request cannot be performed because the disk shelf does not exist.

**Corrective Action**

Run the test against a shelf that exists.

**Syslog Message**

Could not perform TEST request; the disk shelf does not exist.

**Parameters**

(None).



# shm events

## shm.bypassed events

### shm.bypassed.disk.fail

#### Deprecated

Deprecated as of Data ONTAP 9.3 because bypass code has been removed.

#### Severity

ERROR

#### Description

This message occurs when the system recommends that a disk be failed because it has been bypassed.

#### Corrective Action

(None).

#### Syslog Message

shm: Disk %s has been recommended for failure due to being bypassed by the shelf.

#### Parameters

**diskName** (STRING): Name of the disk.

### shm.bypassed.disk.fail.disabled

#### Deprecated

Deprecated as of Data ONTAP 9.3 because bypass code has been removed.

#### Severity

ALERT

#### Description

This message occurs when three or more disks have been bypassed on a port.

#### Corrective Action

(None).

#### Syslog Message

shm: Disk bypass check has been disabled due to multiple bypassed disks on host bus adapter %s, shelf %d.

#### Parameters

**hba** (STRING): Name of the host bus adapter.

**shelf** (INT): Shelf ID.

### shm.bypassed.disk.fail.enabled

**Severity**

NOTICE

**Description**

This message occurs when the bypass disabled state is cleared.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk bypass check has been enabled for host bus adapter %s.

**Parameters**

**hba** (STRING): Name of the host bus adapter.

## shm.disk events

### shm.disk.lostWriteError

**Severity**

ERROR

**Description**

This message occurs when the system detects a lost write error.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has detected and recovered from a lost write error; the system will fail the disk if possible.

**Parameters**

**diskName** (STRING): Name of the disk.

## shm.fab events

### shm.fab.writeError

**Severity**

NOTICE

**Description**

This message occurs when an error is received while writing the failure bytes.

**Corrective Action**

(None).

**Syslog Message**

shm: Unable to write failure bytes to disk %s due to error %x.

## Parameters

**diskName** (STRING): Name of the disk.  
**error\_code** (INTECH): Error code returned.

## shm.fab.writeSenseError

### Severity

NOTICE

### Description

This message occurs when an error is received while writing the failure bytes.

### Corrective Action

(None).

### Syslog Message

shm: Unable to write failure bytes to disk %s due to error %x/%x/%x/%x.

## Parameters

**diskName** (STRING): Name of the disk.  
**sense\_key** (INTECH): Actual sense key.  
**sense\_code** (INTECH): Actual sense code.  
**qualifier** (INTECH): Actual sense code qualifier.  
**fru\_failed** (INTECH): Actual FRU code.

## shm.failurebytes events

### shm.failureBytes.fullError

#### Severity

NOTICE

#### Description

This message occurs when writing the failure byte area returns an error due to being full.

#### Corrective Action

(None).

#### Syslog Message

shm: Disk %s failure area has reached capacity; to clear, reload drive firmware.

#### Parameters

**diskName** (STRING): Name of the disk.

## shm.nvme events

### shm.NVME.critical.warning

**Severity**

ERROR

**Description**

This message occurs when the NVMe drive returns a critical warning to the health monitor.

**Corrective Action**

(None).

**Syslog Message**

Disk %s returned a critical warning 0x%x and will be recommended for failure.

**Parameters**

**disk\_name** (STRING): Name of the disk.

**warning\_code** (INTHEX): Warning code returned.

## shm.pullogwarning events

### shm.pullLogWarning

**Severity**

NOTICE

**Description**

This message occurs when the disk dongle requests the system to retrieve the error log.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has returned %d warnings and the log will be saved.

**Parameters**

**diskName** (STRING): Name of the disk.

**count** (INT): Count of warnings received.

## shm.reformat events

### shm.reformat.diskError

**Severity**

ERROR

**Description**

This message occurs when the system tries to reformat a disk, but an error is returned.

**Corrective Action**

(None).

## Syslog Message

shm: Unable to format disk %s as it is not supported.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.reformat.partnerDisk

### Severity

NOTICE

### Description

This message occurs when the system detects a drive that started formatting during a sanitization, but did not finish.

### Corrective Action

Reboot the partner node to clear the formatting condition. Rerun the original sanitize operation.

## Syslog Message

shm: Disk %s on the partner has possibly not completed format. To reformat the disk, the partner must be rebooted.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.reformat.startMsg

### Severity

NOTICE

### Description

This message occurs when the system has begun to reformat a disk that was not completed during sanitization.

### Corrective Action

(None).

## Syslog Message

shm: Disk %s is being reformatted due to an error during disk sanitize.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.ssd events

### shm.ssd.threshold.ioLatency

### Severity

NOTICE

## Description

This message occurs when ONTAP® detects that a solid-state drive (SSD) has an average block latency during a five-minute timeframe that is greater than expected. ONTAP will continue to monitor the SSD latency and might fail the SSD if it exceeds latency thresholds.

## Corrective Action

(None).

## Syslog Message

SSD %s has exceeded the expected block latency in the current timeframe with an average latency of %d us and an average utilization of %d percent. The next highest SSD latency: %d us. %s

## Parameters

**disk\_name** (STRING): Name of the disk.

**latency\_avg** (INT): Average block latency in the current timeframe in microseconds (us).

**util\_avg** (INT): Average drive utilization percent in the current timeframe.

**next\_hi\_latency\_avg** (INT): Average block latency for the SSD with the next highest block latency in the current timeframe.

**disk\_information** (STRING): Formatted information about the disk. This includes the disk's vendor, model, firmware revision, and serial number.

# shm.threshold events

## shm.threshold.agrsvIOCount

### Deprecated

This event is removed as it is a test message not intended for customers.

### Severity

NOTICE

### Description

This message occurs when the system detects a disk that exceeds the threshold count of high aggressive IO latencies.

### Corrective Action

(None).

### Syslog Message

Disk %s has exceeded %d IOs which have latencies greater than the threshold.

### Parameters

**disk\_name** (STRING): Name of the disk.

**agrsv\_IO\_count** (INT): Count of aggressive timeout IOs.

## shm.threshold.allMediaErrors

### Severity

ERROR

**Description**

This message occurs when the system detect more than 25 medium and/or recovered errors in a 10-minute window.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has crossed the combination media error threshold in a 10 minute window.

**Parameters**

**diskName** (STRING): Name of the disk.

**shm.threshold.checksumErrors****Severity**

ERROR

**Description**

This message occurs when a disk exceeds the threshold of checksum errors on the same block.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has exceeded the threshold for checksum errors on the same block; the system will fail the disk.

**Parameters**

**diskName** (STRING): Name of the disk.

**shm.threshold.consecutiveAborts****Severity**

ERROR

**Description**

This message occurs when a disk exceeds the threshold of consecutive abort errors on one drive.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has exceeded the threshold of %d consecutive abort errors; the system will fail the disk if possible.

**Parameters**

**diskName** (STRING): Name of the disk.

**count** (INT): Count of errors.

## shm.threshold.consecutiveTimeouts

### Severity

ERROR

### Description

This message occurs when a disk exceeds the threshold of consecutive timeouts on one drive.

### Corrective Action

(None).

### Syslog Message

shm: Disk %s has exceeded the threshold of %d consecutive timeouts; the system will fail the disk if possible.

### Parameters

**diskName** (STRING): Name of the disk.

**count** (INT): Count of errors.

## shm.threshold.disk.pcycle

### Severity

ERROR

### Description

This message occurs when a disk exceeds the threshold of power-cycle error recovery tries.

### Corrective Action

(None).

### Syslog Message

shm: Disk %s has exceeded the threshold for power-cycle error recovery events; the system will fail the disk if possible.

### Parameters

**diskName** (STRING): Name of the disk.

## shm.threshold.disk.poh

### Severity

NOTICE

### Description

This message occurs when the storage system contains drives that exceed five years of cumulative power-on hours. This five-year period represents the expected reliable lifespan of the drive. The likelihood that a drive will fail increases significantly the longer the drive stays in use beyond its warranty period.

### Corrective Action

For additional guidance, search for Support Bulletin SU464 on the NetApp Support Site.



## Syslog Message

There are "%d" drives that have exceeded five years of life. To view them, use the "storage disk show -power-on-hours >=43800 -fields power-on-hours" command.

## Parameters

**count** (INT): Count of drives that have exceeded the five year threshold.

## shm.threshold.highIOLatency

### Severity

ERROR

### Description

This message occurs when the system detects a disk that has an average IO latency during the current window that is significantly greater than all other drives of the same class.

### Corrective Action

(None).

## Syslog Message

Disk %s exceeds the average IO latency threshold and will be recommended for failure.

## Parameters

**disk\_name** (STRING): Name of the disk.

## shm.threshold.lipStormReset

### Severity

ERROR

### Description

This message occurs when the system detects more than one instance of a disk resetting itself because of a large number of Loop Initialization Procedure requests (a LIP storm). The disk will be failed.

### Corrective Action

Replace the disk.

## Syslog Message

shm: The system has detected more than one LIP storm reset on disk %s.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.threshold.mediaErrorsLba

### Severity

ERROR

### Description

This message occurs when the system detects more than three media (recovered or medium) errors on the same block.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has had multiple media errors on sector %llu.

**Parameters**

**diskName** (STRING): Name of the disk.

**block\_num** (LONGINT): Block number.

**shm.threshold.mediaErrorsReassign****Severity**

ERROR

**Description**

This message occurs when the system detects more than three medium and/or recovered errors in a 10-minute window on the same disk sector. The bad sector will be reassigned.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has had multiple media errors on the same sector in the last 10 minutes, and is reassigning the sector.

**Parameters**

**diskName** (STRING): Name of the disk.

**shm.threshold.mediumErrors****Severity**

ERROR

**Description**

This message occurs when the system detects more than 25 medium errors in a 10-minute window.

**Corrective Action**

(None).

**Syslog Message**

shm: Disk %s has crossed the medium error threshold in a 10 minute window.

**Parameters**

**diskName** (STRING): Name of the disk.

**shm.threshold.mediumErrors7days****Severity**

ERROR

## Description

This message occurs when the system detects more than 100 medium errors in a seven-day window.

## Corrective Action

(None).

## Syslog Message

shm: Disk %s has crossed the medium error threshold in a seven-day window.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.threshold.ratedLife

### Severity

NOTICE

## Description

This message occurs when the rated life used by a solid state drive (SSD) exceeds 90%. When an SSD reaches 100% of rated life, it might not be able to retain data while powered off for long periods of time.

## Corrective Action

The number of weeks of rated life remaining indicated for each SSD is an estimate based on past usage. If the SSD is expected to remain in service beyond the estimated remaining time, initiate planning for SSD replacement when the rated life reaches 100%. Use the "storage disk show -ssd-wear" command to monitor the current rated life used by your SSDs.

## Syslog Message

shm: There are %d drives that have consumed at least 90 percent of their rated life: %s.

## Parameters

**count** (INT): Count of drives that have exceeded the threshold.

**disk\_names** (STRING): Name, percentage of rated life and estimated remaining time before reaching 100% rated life of each SSD that has exceeded the threshold.

## shm.threshold.ratedLife2

### Severity

ERROR

## Description

This message occurs when the rated life used by a solid state drive (SSD) exceeds 95%. When an SSD reaches 100% of rated life, it might not be able to retain data while powered off for long periods of time.

## Corrective Action

The number of weeks of rated life remaining indicated for each SSD is an estimate based on past usage. If the SSD is expected to remain in service beyond the estimated remaining time, initiate planning for SSD replacement when the rated life reaches 100%. Use the "storage disk show -ssd-wear" command to monitor the current rated life used by your SSDs.

## Syslog Message

shm: There are %d drives that have consumed at least 95 percent of their rated life: %s.

## Parameters

**count** (INT): Count of drives that have exceeded the threshold.

**disk\_names** (STRING): Name, percentage of rated life and estimated remaining time before reaching 100% rated life of each SSD that has exceeded the threshold.

## shm.threshold.ratedLifeMax

### Severity

ALERT

### Description

This message occurs when the rated life used value exceeds 100% on a solid state drive (SSD). When an SSD reaches 100% of rated life, it might not be able to retain data while powered off for long periods of time.

### Corrective Action

Replace the SSDs that have reached the end of their rated life.

## Syslog Message

shm: There are %d drives that have reached the end of their rated life: %s

## Parameters

**count** (INT): Count of drives that have exceeded the threshold.

**disk\_names** (STRING): Name and percentage of rated life of each SSD that has exceeded the threshold.

## shm.threshold.recoveredErrors

### Severity

ERROR

### Description

This message occurs when the system detects more than 25 recovered errors in a 10-minute window.

### Corrective Action

(None).

## Syslog Message

shm: Disk %s has crossed the recovered error threshold in a 10 minute window.

## Parameters

**diskName** (STRING): Name of the disk.

## shm.threshold.sensekey

### Severity

ERROR

## Description

This message occurs when a disk exceeds the threshold of a particular sense key error.

## Corrective Action

(None).

## Syslog Message

shm: Disk %s has exceeded the threshold for sense key %d errors; the system will fail the disk if possible.

## Parameters

**diskName** (STRING): Name of the disk.

**senseKey** (INT): Sense key value.

## shm.threshold.spareBlocksConsumed

### Severity

NOTICE

### Description

This message occurs when the spare blocks consumed value exceeds the first threshold on an SSD drive.

### Corrective Action

(None).

### Syslog Message

shm: There are %d disks that have consumed at least 60 percent of their use-based internal spare capacity. The affected disks are: %s.

### Parameters

**count** (INT): Count of disks which have exceeded the threshold.

**disk\_names** (STRING): Names of the disk drives.

## shm.threshold.spareBlocksConsumedMax

### Severity

NOTICE

### Description

This message occurs when the spare blocks consumed value exceeds the second threshold on an SSD drive.

### Corrective Action

(None).

### Syslog Message

shm: There are %d disks that have consumed at least 80 percent of their use-based internal spare capacity. The affected disks are: %s.

### Parameters

**count** (INT): Count of disks which have exceeded the second threshold.

**disk\_names** (STRING): Names of the disk drives.

# showmount events

## showmount.update events

### showmount.update.failed

**Severity**

ERROR

**Description**

This message occurs when the showmount job fails to update the "exports.data" file. This file provides a list of exported volumes and qtrees on the NFS server to any clients running the "showmount -e" command.

**Corrective Action**

Verify that the root volume of the Vserver is online and not full. After the Vserver is online and space is available, run the following command: "export-policy cache flush -vserver <vserver-name> -cache showmount".

**Syslog Message**

Flushing to exports.data file failed on Vserver "%s".

**Parameters**

**vserver** (STRING): Name of the Vserver where the problem is occurring.

# sis events

## sis.aggr events

### sis.aggr.cvs.undo.abort

#### Severity

NOTICE

#### Description

This message occurs when undo of cross volume savings on aggregate has to abort due to some error conditions.

#### Corrective Action

(None).

#### Syslog Message

Undo of cross volume savings on aggregate %s aborted. Volume %s encountered error: %s

#### Parameters

**aggr\_name** (STRING): Aggregate name.  
**vvol\_name** (STRING): Flex volume name.  
**error** (STRING): Error code for the abort.

### sis.aggr.metadir.create.failed

#### Severity

ERROR

#### Description

This message occurs when the creation of SIS metafiles in the aggregate containing the SIS volume fails. Change logging will be turned off on the volume.

#### Corrective Action

Ensure that space is available in the aggregate containing the volume. Use the 'sis on' command to turn on change logging on the volume, which attempts to re-create SIS metafiles in the aggregate. If enough space is available, start wafliion on the aggregate containing the indicated flexible volume.

#### Syslog Message

Creation of aggregate metadir failed for SIS volume %s%s. Change logging will be turned off on the volume: %s.

#### Parameters

**volname** (STRING): SIS volume name.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.  
**reason** (STRING): Reason for the failure.

## sis.auto events

### sis.auto.session.change

#### Severity

NOTICE

#### Description

This message occurs whenever the number of background deduplication sessions change on a node because of throttling based on system load for the volumes with efficiency policy type as auto.

#### Corrective Action

(None).

#### Syslog Message

ADS: Number of auto sessions changed from %d to %d

#### Parameters

**old\_sessions** (INT): prev sessions count.

**new\_sessions** (INT): new sessions count.

## sis.autosched events

### sis.autoSched.failed

#### Severity

ERROR

#### Description

This event is issued when a SIS volume operation could not be auto scheduled due to unexpected conditions.

#### Corrective Action

(None).

#### Syslog Message

SIS operation for %s%s could not be auto scheduled: %s

#### Parameters

**vol** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Error message

## sis.c2c events

### sis.c2c.start.fail

#### Severity

ERROR



## Description

This event is issued when dense scanner could not start on c2c restore volume during restore operation.

## Corrective Action

Run "sis start" command to manually start the SIS scanner.

## Syslog Message

SIS scanner on C2C restore volume %s%s could not be started: %s.

## Parameters

**vol** (STRING): The name of the volume

**valident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Error message

# sis.cfg events

## sis.cfg.deleted

### Severity

ERROR

### Description

This message occurs when the configuration file for a Single Instance Storage (SIS) volume is deleted. This leads to the loss of deduplication configuration for the volume and statistics of deduplication operations.

### Corrective Action

Remount the volume, enable deduplication using the "sis on" command, configure the volume using the "sis config" command, and then run the "sis start -s -m" command to start building deduplication metadata.

### Syslog Message

SIS configuration for volume %s%s was deleted: %s. Remount the volume, then enable SIS and reconfigure the volume to resume deduplication.

### Parameters

**path** (STRING): Path to a SIS volume.

**valident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Reason for deletion.

## sis.cfg.invalidEntry

### Severity

ERROR

### Description

This event is issued when a SIS volume operation found an invalid configuration for the volume.

### Corrective Action

Automatically checks and creates default configuration after detection. Please use 'sis' command to change default settings.

## Syslog Message

SIS volume configuration: found invalid SIS configuration file %s%s: %s

## Parameters

**path** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**reason** (STRING): Reason why SIS operation considered it as wrong.

## sis.cfg.setFailed

### Severity

ERROR

### Description

This message occurs when a Single Instance Storage (SIS) volume operation fails to write in the configuration file. This might lead to incorrect reporting of statistics. If the error occurred during a Data ONTAP® version upgrade, it might also lead to the upgraded configuration not being loaded. The volume might not appear in the 'sis status' command output as a result. Also, the system might not allow changing the configuration for that volume.

### Corrective Action

Ensure that space is available in the indicated volume. If this error occurred during a Data ONTAP version upgrade and the volume does not appear in the 'sis status' command output, remount the volume after correcting the underlying condition. If enough space is available, start wafiron on the aggregate containing the indicated flexible volume.

## Syslog Message

Saving SIS volume configuration for volume %s%s : %s

## Parameters

**path** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**reason** (STRING): Reason for the failure.

## sis.changelog events

### sis.changelog.full

#### Severity

ERROR

#### Description

This message occurs when the Single Instance Storage (SIS) change logging metafile is full. Change logging is stopped. New writes to the volume will not be change logged until the next SIS operation.

#### Corrective Action

Issue the 'sis start' command on the volume to process existing changelogs. To process changelog files more frequently before they reach full, adjust the SIS operation schedule with the 'sis config -s' command. If there is significant loss of savings, then run the 'sis start -s' command, allowing it to run to completion.

## Syslog Message

SIS change logging metafile for volume %s%s%s is full.

## Parameters

**host** (STRING): Owing host of the volume.

**vol** (STRING): Name of the volume.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

## sis.changelog.noinit

### Severity

ERROR

### Description

This event is issued when the change logging operation could not be initialized on a SIS volume.

### Corrective Action

Run "sis on" command to reinitialzie the change logging operation.

## Syslog Message

Change logging operation on volume %s%s could not be initialized: %s.

## Parameters

**vol** (STRING): The name of the volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Error message

## sis.changelog.noload

### Severity

NOTICE

### Description

This event is issued when prior change logging metafiles could not be reused during the initialization. New metafiles will be created for the operation.

### Corrective Action

(None).

## Syslog Message

Prior Change logging metafiles on volume %s%s could not be reused.

## Parameters

**vol** (STRING): The name of the volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.changelog.truncate

### Severity

NOTICE

## Description

This message occurs when the Single Instance Storage (SIS) change logging metafile is truncated as part of ADS session.

## Corrective Action

(None).

## Syslog Message

Prior Change logging metafiles on volume %s%s are truncated.

## Parameters

**host** (STRING): Owing host of the volume.

**vol** (STRING): Name of the volume.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

# sis.chkpoint events

## sis.chkpoint.delete

### Severity

INFORMATIONAL

### Description

This message occurs when a Single Instance Storage (SIS) checkpoint is deleted. The current SIS operation is restarted from the beginning.

### Corrective Action

(None).

### Syslog Message

SIS checkpoint on the volume '%s%s' is deleted. Starting current SIS operation from the beginning.

### Parameters

**volume** (STRING): Volume path.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.chkpoint.restore.failed

### Severity

ERROR

### Description

This message occurs when a Single Instance Storage (SIS) operation fails to restore from a previous checkpoint. The SIS operation is restarted from the beginning.

### Corrective Action

(None).

### Syslog Message

SIS operation on the volume '%s%s' failed to restore from a previous checkpoint. Starting SIS operation from the beginning.

## Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.clog events

## sis.clog.disabled

### Severity

ERROR

### Description

This message occurs while mounting a volume, when an error with the SIS metafile operation results in the change logging function being disabled to avoid inconsistencies.

### Corrective Action

Reinitialize the change logging operation by using the 'sis on' command.

### Syslog Message

Change logging has been disabled for %s%s due to an error that occurred while mounting the volume during the SIS metafile operation: %s.

## Parameters

**path** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Reason for the failure.

# sis.compress events

## sis.compress.skip

### Severity

INFORMATIONAL

### Description

This message occurs when postprocess compression is skipped during a deduplication operation, even though postprocess compression is enabled on the volume. Post process compression is not supported on a node configured with the All-Flash Optimized personality.

### Corrective Action

(None).

### Syslog Message

Skipped postprocess compression on volume %s%s because the node is All-Flash Optimized.

## Parameters

**path** (STRING): Path to the deduplicated volume.

**volident** (STRING): Unique volume identifier.

# sis.corrupt events

## sis.corrupt.metafile.deleted

### Severity

NOTICE

### Description

This message occurs when a Single Instance Storage (SIS) metafile is deleted because there are Disk/Media errors affecting the file.

### Corrective Action

The deleted metafile might be recovered automatically the next time you run the 'sis start' command. If you have not previously run the 'sis start' command on this volume, run the 'sis start -s' command to generate metafiles and obtain efficient deduplication. Contact NetApp technical support for guidance.

### Syslog Message

SIS metafile on volume '%s%s' is deleted.

### Parameters

**volume** (STRING): Volume path.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.default events

## sis.default.auto.se.enable

### Severity

NOTICE

### Description

This message occurs whenever volume move is done from a non All Flash Optimized node to an All Flash Optimized node and we attempt to enable storage efficiency features and automatic offline volume efficiency.

### Corrective Action

(None).

### Syslog Message

Volume efficiency features will be enabled for volume %s on aggregate %s.

### Parameters

**vvol\_name** (STRING): Flex volume name.

**aggr\_name** (STRING): Aggregate name.

## sis.default.auto.se.enable.fail

### Severity

NOTICE

## Description

This message occurs whenever volume move is done from a non All Flash Optimized node to an All Flash Optimized node and the attempt to enable automatic offline volume efficiency failed.

## Corrective Action

(None).

## Syslog Message

Failed to enable automatic offline volume efficiency features for volume %s on aggregate %s.

## Parameters

**vvol\_name** (STRING): Flex volume name.

**aggr\_name** (STRING): Aggregate name.

**reason** (STRING): Reason for the failure.

# sis.enable events

## sis.enable.post.sm.break

### Severity

NOTICE

### Description

This message occurs when the "storage efficiency" feature is enabled on a destination volume due to a broken SnapMirror® relationship on an All Flash Optimized system. When the relationship breaks, the DP volume is converted to RW, and all RW volumes have the "storage efficiency" feature enabled.

### Corrective Action

(None).

### Syslog Message

"Storage efficiency" feature has been enabled on the destination volume %s due to a SnapMirror break operation.

### Parameters

**volume** (STRING): Destination FlexVol® volume name on which the "storage efficiency" feature is enabled.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.enable.post.ssm.init

### Severity

NOTICE

### Description

This message occurs when the "storage efficiency" feature is enabled on a destination volume with Synchronous SnapMirror® on an All Flash Optimized system.

### Corrective Action

(None).

## Syslog Message

"Storage efficiency" feature has been enabled on the destination volume %s after Asynchronous Snapmirror converted to Synchronous Snapmirror.

## Parameters

**volume** (STRING): Destination FlexVol® volume name on which the "storage efficiency" feature is enabled.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.fingerprint events

## sis.fingerprint.notify.lost

### Severity

NOTICE

### Description

This message occurs when a bad block is encountered while reading SIS metafiles. As a result, some of the fingerprint records used to identify blocks for use in deduplication might have been lost. The only impact from this is lower efficiency in deduplication.

### Corrective Action

(None).

## Syslog Message

Bad block found while reading SIS metafiles for volume %s%. As a result some fingerprint data might have been lost and could result in lower space savings.

## Parameters

**volume** (STRING): Volume name.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.idedup events

## sis.idedup.disable

### Severity

NOTICE

### Description

This message occurs when inline deduplication is disabled if aggregate is not hybrid or not part of All-Flash Optimized node since inline deduplication is not supported on them.

### Corrective Action

(None).

## Syslog Message

Disabled inline deduplication on volume %s% because the aggregate is not hybrid or not part of All-Flash Optimized node.



## Parameters

**path** (STRING): Path to the deduplicated volume.

**volident** (STRING): Unique volume identifier.

# sis.incorrect events

## sis.incorrect.duration.notify

### Severity

ERROR

### Description

This message occurs when a Single Instance Storage (SIS) volume repeatedly can not finish SIS operation within user-given 'duration' settings.

### Corrective Action

Change duration settings appropriately, based on SIS completion time history and total amount data that needs to be processed.

### Syslog Message

SIS operation repeatedly can not finish within duration limits specified by user.

## Parameters

**volname** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.logical events

## sis.logical.limit.near

### Severity

NOTICE

### Description

This message occurs when the deduplication engine's logical data limit is nearly reached on a volume. Logical data is the summation of used blocks, saved blocks, and compressed blocks in the volume.

### Corrective Action

Move data to another volume or delete unwanted data from the volume.

### Syslog Message

Deduplication engine's logical data limit is nearly reached on volume %s%s.

## Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.logical.limit.reached

### Severity

NOTICE

### Description

This message occurs when the deduplication engine's logical data limit is reached on the volume. Deduplication operations might not run on the volume. Logical data is the summation of used blocks, saved blocks, and compressed blocks in the volume.

### Corrective Action

Move data to another volume or delete unwanted data from the volume.

### Syslog Message

Deduplication engine's logical data limit was reached on volume %s%s.

### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.max events

### sis.max.saving.near

#### Severity

NOTICE

#### Description

This message occurs when Single Instance Storage (SIS) savings on the volume is nearing its maximum limit. No further savings can be realized on the volume after the limit is reached.

#### Corrective Action

(None).

#### Syslog Message

Maximum savings limit nearly reached on volume %s%s.

#### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

### sis.max.saving.reached

#### Severity

NOTICE

#### Description

This message occurs when Single Instance Storage (SIS) savings on the volume reaches its maximum limit. No further savings can be realized on the volume.

## Corrective Action

(None).

## Syslog Message

Maximum savings limit reached on volume %s%s.

## Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.metadata events

## sis.metadata.inconsistent

### Severity

ERROR

### Description

This message occurs when inconsistencies are found in deduplication metadata. This might be due to a failed metadata upgrade. The impact will be very low deduplication savings.

### Corrective Action

Run the 'sis start -s -m -vserver vserver\_name -volume volume\_name' command. This rebuilds the deduplication metadata. If the command fails or if subsequent deduplication runs do not provide adequate space savings, contact NetApp technical support.

### Syslog Message

Deduplication metadata is inconsistent in volume: %s%s.

### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): Unique volume identifier for cases when the volume name alone is insufficient.

# sis.op events

## sis.op.aborted

### Severity

ERROR

### Description

This event is issued when a SIS volume operation is aborted due to unexpected conditions.

### Corrective Action

(None).

### Syslog Message

SIS operation for %s%s has aborted: %s

## Parameters

**path** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Error message

## sis.op.idle

### Severity

NOTICE

### Description

This message occurs when deduplication is enabled on the volume and deduplication process has not been run for more than configured idle threshold time set(default 7 days).

### Corrective Action

Run the "sis start" command.

### Syslog Message

SIS operation is idle on volume "%s" since %d weeks %d days.

## Parameters

**volpath** (STRING): Volume name.

**last\_dedup\_time** (INT): Last Successful Deduplication operation on volume in weeks.

**last\_dedup\_time\_days** (INT): Last Successful Deduplication operation on volume in days.

## sis.op.stopped

### Severity

INFORMATIONAL

### Description

This event is issued when a SIS volume operation is stopped by user.

### Corrective Action

(None).

### Syslog Message

SIS operation for %s%s has stopped

## Parameters

**path** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.policy events

### sis.policy.predef.create.upgradeError

### Severity

NOTICE

## Description

This message occurs during system software upgrade to Data ONTAP® 8.1 or later, while creating predefined policies. There is a user policy with the same name as a predefined policy prior to upgrade, and it cannot be renamed. Therefore, the predefined policy cannot be created.

## Corrective Action

Rename pre-existing user policy and run system software upgrade again.

## Syslog Message

Policy upgrade failed because the user policy %s could not be renamed to %s.

## Parameters

**old\_user\_policy** (STRING): Name of the preexisting user policy.

**renamed\_user\_policy** (STRING): New name for the preexisting user policy.

## sis.policy.rename.upgrade

### Severity

INFORMATIONAL

### Description

This message occurs during upgrade to Data ONTAP® 8.1 or later, while creating predefined efficiency policies. The user policy was renamed because of the creation of predefined policies.

### Corrective Action

(None).

### Syslog Message

During upgrade, user policy name %s is changed to %s.

### Parameters

**old\_user\_policy** (STRING): Name of the preexisting user policy.

**renamed\_user\_policy** (STRING): New name for the preexisting user policy.

## sis.reg events

### sis.reg.invalidEntry

#### Severity

ERROR

#### Description

This event is issued when a SIS volume operation found an invalid registry entry.

#### Corrective Action

(None).

#### Syslog Message

SIS volume registry: found invalid SIS registry entry for %s%s: %s

## Parameters

**path** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**regval** (STRING): Registry value

## sis.revert events

### sis.revert.to.has.Aborted

#### Severity

NOTICE

#### Description

This message occurs when the Single Instance Storage (SIS) volume command 'sis revert\_to' has failed. This might prevent running generic Data ONTAP® 'revert\_to' command successfully

#### Corrective Action

(None).

#### Syslog Message

sis 'revert\_to' command has failed on volume %s%s :%s

## Parameters

**path** (STRING): Path to the SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**reason** (STRING): Reason for the failure.

### sis.revert.to.has.Succeeded

#### Severity

NOTICE

#### Description

This message occurs when the Single Instance Storage (SIS) volume command 'sis revert\_to' has succeeded.

#### Corrective Action

(None).

#### Syslog Message

'sis revert\_to' command has succeeded on volume %s%s

## Parameters

**path** (STRING): Path to the SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

### sis.revert.to.large.vol

#### Severity

ERROR

## Description

This message occurs when the Single Instance Storage (SIS) volume command 'sis revert\_to' fails because the logical data exceeded the maximum allowed for SIS in the target release. This is to prevent growing deduplication metadata beyond the size that is supported in the target release. Logical data is the summation of used blocks, saved blocks, and compressed blocks in the volume. This might prevent running the generic Data ONTAP® 'revert\_to' command successfully.

## Corrective Action

Shrink the volume or run 'sis revert\_to -delete vol\_path' to delete deduplication metadata from this volume.

## Syslog Message

'sis revert\_to' command has failed on volume %s%s. Volume exceeded the maximum allowed logical data limit of %d TB for SIS in the target release.

## Parameters

**path** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**maxVolSize** (INT): Maximum size of the SIS volume in the target release.

## sis.revert.to.space.warning

### Severity

INFORMATIONAL

### Description

This message occurs when a Single Instance Storage (SIS) volume command 'sis revert\_to' is run. This message indicates that adding more data during this command execution might abort that operation due to insufficient space.

### Corrective Action

(None).

### Syslog Message

Adding data to volume %s%s while the 'sis revert\_to' command is running might lead to insufficient free space and failure of that operation.

### Parameters

**path** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## sis.revert.to.write.Failed

### Severity

ERROR

### Description

This message occurs when the Single Instance Storage (SIS) volume command 'sis revert\_to' fails to write to a file or create a file related to a revert\_to operation. This might prevent running generic Data ONTAP® 'revert\_to' command successfully

## Corrective Action

Ensure that space is available in the indicated volume or in the aggregate it belongs to. After space is available, rerun the 'sis revert\_to' command.

## Syslog Message

'sis revert\_to' command failed on volume %s%s :%s.

## Parameters

**path** (STRING): Path to the SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**reason** (STRING): Reason for the failure.

# sis.revertmetacreate events

## sis.revertmetacreate.failed

### Severity

ERROR

### Description

This message occurs when Single Instance Storage (SIS) metafiles cannot be created in the flexible volume metadir during a reversion to an earlier version of Data ONTAP®.

### Corrective Action

Following this error, SIS is disabled for the volume. After correcting the underlying issue indicated by the error message, enable SIS again using the 'sis on' command. SIS then resumes normally.

### Syslog Message

Could not create the metafiles for the volume %s%s in the volume metadir: %s SIS is turned off for this volume.

### Parameters

**vol** (STRING): Path to the volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Reason for the failure.

# sis.schedule events

## sis.schedule.notify.change

### Severity

NOTICE

### Description

This message occurs when the Single Instance Storage (SIS) schedule of a volume is changed to the default schedule (sun-san@0) because the node's clusterid does not match with the SIS config clusterid. This can happen when a volume is replicated to a node belonging to another cluster.



### Corrective Action

Use the `sis config` command to set a SIS policy on the volume.

### Syslog Message

SIS policy for volume %s%s is no longer valid because of the clusterid mismatch. SIS schedule is set to the default schedule `sun-sat@0`.

### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.schedule.notify.change.compression

### Severity

NOTICE

### Description

This message occurs when a volume with background compression has an auto type policy. The policy is replaced with a default schedule (`sun-sat@0`). This can happen when a volume is being re-mounted during which efficiency policy type is changed to auto.

### Corrective Action

Use the `volume efficiency modify` command to set a efficiency policy on the volume. For example: `volume efficiency modify -volume volume-name -policy policy-name`

### Syslog Message

Efficiency policy for volume %s%s is auto type and is not supported with background compression. The schedule is set to the default schedule `sun-sat@0`.

### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.schedule.notify.change.snapmirror

### Severity

NOTICE

### Description

This message occurs when a SnapMirror® relationship spanning across Vservers breaks. In such cases, the efficiency policy on the destination is no longer valid; it is replaced with the default schedule (`sun-sat@0`).

### Corrective Action

Use the `'volume efficiency modify'` command to set a efficiency policy on the volume. For example: `volume efficiency modify -volume volume-name -policy policy-name`

### Syslog Message

Efficiency policy for volume %s%s is not valid and has been replaced with default schedule (`sun-sat@0`).

## Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.undo events

## sis.undo.aborted

### Severity

ERROR

### Description

This messages occurs when a SIS volume 'undo' operation is aborted due to unexpected conditions.

### Corrective Action

(None).

### Syslog Message

SIS undo operation for '%s%s' has aborted: '%s'.

## Parameters

**path** (STRING): Path to a SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Error message.

## sis.undo.aborted.inconsistent

### Severity

ALERT

### Description

This message occurs when a SIS volume 'undo' operation is aborted due to volume inconsistency.

### Corrective Action

Run waflliron to remove inconsistencies in the volume, then retry the operation.

### Syslog Message

SIS 'undo' operation for '%s%s' has aborted due to volume inconsistencies.

## Parameters

**path** (STRING): Path to a SIS volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.undo.done

### Severity

NOTICE

## Description

This message occurs when 'sis undo -t wrong' command removes the sharing on improperly shared inodes.

## Corrective Action

(None).

## Syslog Message

The 'sis undo -t wrong' command removed sharings on improperly shared inodes in the volume '%s%s'. Details about cleaned inodes are available at '%s'.

## Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**log\_location** (STRING): SIS undo log location.

## sis.undo.nospace

### Severity

ERROR

### Description

This event is issued to indicate that there is not enough free space to undo the shared blocks on a SIS volume. The undo operation will be aborted. Data blocks that were still shared at the time the undo operation stopped will remain shared. Blocks already processed by the undo operation are no longer shared.

### Corrective Action

Free disk space in the volume, or increase the size of the volume, then retry the operation.

### Syslog Message

Undoing shared blocks on volume %s%s has aborted because there is insufficient free space. Data blocks that were still shared at the time the undo operation stopped will remain shared. Blocks already processed by the undo operation are no longer shared.

### Parameters

**vol** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.undo.required

### Severity

NOTICE

### Description

This message occurs when 'sis undo -t wrong -d' command finds improperly shared inodes.

### Corrective Action

Run the 'sis undo -t wrong' or 'sis undo -i' command to remove improperly shared inodes.

### Syslog Message

The 'sis undo -t wrong -d' command detected improperly shared inodes in the volume '%s%s'. Run the 'sis

undo -t wrong' or 'sis undo -i' command to remove improperly shared inodes. Details about improperly shared inodes are available at '%s'.

### Parameters

**volume** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**log\_location** (STRING): SIS undo log location.

## sis.undo.started

### Severity

INFORMATIONAL

### Description

This message occurs when a SIS undo operation is started.

### Corrective Action

(None).

### Syslog Message

SIS undo operation for %s%s has started.

### Parameters

**path** (STRING): Path to a SIS volume.

**volident** (STRING): Uniquely identifies the volume when the volume name alone is insufficient.

## sis.undo.stopped

### Severity

ERROR

### Description

This event is issued when a SIS volume undo operation is stopped by user.

### Corrective Action

(None).

### Syslog Message

SIS undo operation for %s%s has stopped.

### Parameters

**path** (STRING): Path to a SIS volume

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## sis.upgrade events

### sis.upgrade.aggr.metfile.failed

### Severity

ERROR

## Description

This message occurs when Single Instance Storage (SIS) metafiles cannot be created in the aggregate containing the volume, during a Data ONTAP® upgrade. The aggregate might be full. The 'sis status' command will not display this volume. SIS operations cannot be triggered manually or according to a schedule for the volume.

## Corrective Action

Ensure that space is available in the aggregate containing the volume. After correcting the underlying issue indicated by the error message, remount the volume for the SIS upgrade to be triggered again. If the upgrade succeeds, SIS then resumes normally.

## Syslog Message

Could not create the metafiles for the volume %s%s : %s.

## Parameters

**volname** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Reason for the failure.

## sis.upgrade.feature.enable

### Severity

NOTICE

### Description

This message occurs during upgrade to Data ONTAP® 9.2 or later, when we start enabling Volume Efficiency features on the existing compatible volumes on All-Flash Optimized nodes in the background.

### Corrective Action

(None).

### Syslog Message

Volume efficiency features will be enabled on Volumes of All-Flash Optimized nodes in the Cluster.

### Parameters

(None).

## sis.upgrade.feature.enable.done

### Severity

NOTICE

### Description

This message occurs during upgrade to Data ONTAP® 9.2 or later, when we have completed enabling the Volume Efficiency features on the existing compatible volumes.

### Corrective Action

(None).

### Syslog Message

Volume efficiency features were enabled on Volumes of All-Flash Optimized nodes in the Cluster.

## Parameters

(None).

## sis.upgrade.feature.enable.fail

### Severity

INFORMATIONAL

### Description

This message occurs during upgrade to Data ONTAP® 9.2 or later, when we fail to enable Volume Efficiency features on the existing compatible volume on All-Flash Optimized nodes.

### Corrective Action

Manually try enabling the Volume Efficiency features on the volume using the "volume efficiency modify" command.

### Syslog Message

Could not enable Volume Efficiency features for Volume %s of vserver %s.%s

### Parameters

**volume** (STRING): Volume name.

**vserver** (STRING): Vserver name.

**reason** (STRING): Reason for failure.

## sis.upgrade.feature.enabled

### Severity

INFORMATIONAL

### Description

This message occurs during upgrade to Data ONTAP® 9.2 or later, when we enable Volume Efficiency features on the existing compatible volume on All-Flash Optimized nodes.

### Corrective Action

(None).

### Syslog Message

Volume Efficiency features enabled on Volume %s of Vserver %s.

### Parameters

**volume** (STRING): Volume name.

**vserver** (STRING): Vserver name.

## sis.upgrademetatile events

### sis.upgrademetatile.failed

### Severity

ERROR

## Description

This message occurs when Single Instance Storage (SIS) metafiles cannot be created during a Data ONTAP® upgrade. The 'sis status' command will not display this volume. SIS operations cannot be triggered manually or according to a schedule for the volume.

## Corrective Action

Correct the underlying issue indicated by the error message, and then run 'sis on' command to trigger the SIS upgrade again. If the upgrade succeeds, SIS then resumes normally.

## Syslog Message

Could not create the metafiles for the volume %s%s: %s.

## Parameters

**vol** (STRING): Path to the volume.

**valident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**err** (STRING): Reason for the failure.

# sis.verify events

## sis.verify.suspended

### Severity

NOTICE

### Description

This message occurs when the system suspends a Single Instance Storage (SIS) volume checking operation to to run a queued SIS operation.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**path** (STRING): Path to the SIS volume.

**valident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# sis.vol events

## sis.vol.auto.promote

### Severity

NOTICE

### Description

This message occurs when deprioritized volume gets auto promoted.

**Corrective Action**

(None).

**Syslog Message**

Volume "%s" is auto promoted: Percentage of writes happened in this window (%d%%) and percentage of dedupe savings (%d%%).

**Parameters**

**vvol\_name** (STRING): Volume name.

**inline\_process\_pct** (INT): Percentage of writes happened after one window.

**inline\_saving\_pct** (INT): Percentage of inline dedupe savings in each window.

**sis.vol.overwrite.depri****Severity**

NOTICE

**Description**

This message occurs when automatic deduplication deprioritizes a volume due to a large number of overwrites in the block fingerprint log file.

**Corrective Action**

(None).

**Syslog Message**

Automatic deduplication on volume "%s" is deprioritized: percentage of overwrites in the fingerprint log file is %d%%.

**Parameters**

**vvol\_name** (STRING): Volume name.

**stale\_pct** (INT): Percentage of cumulative stale fingerprints after the fifth iteration of automatic deduplication runs.

**sis.vol.savings.depri****Severity**

NOTICE

**Description**

This message occurs when automatic deduplication deprioritizes a volume due to potentially low storage efficiency savings.

**Corrective Action**

(None).

**Syslog Message**

Automatic deduplication on volume "%s" is deprioritized: low background savings (%d%%), low cross volume deduplication savings (%d%%) and low inline savings (%d%%) when compared total with minimum required (%d%%) .



## Parameters

**vvol\_name** (STRING): Volume name.

**bkg\_savings** (INT): Background deduplication savings realized after the fifth iteration of automatic deduplication runs.

**cvbd\_savings** (INT): Cross Volume deduplication savings realized on the volume.

**inline\_savings** (INT): Inline deduplication savings realized on the volume.

**min\_savings\_req** (INT): Min Savings required

## sis.vol.verify.depri

### Severity

NOTICE

### Description

This message occurs when automatic deduplication deprioritizes a volume due to a large number of stale fingerprint removal operations on the volume.

### Corrective Action

(None).

### Syslog Message

Automatic deduplication on volume "%s" is deprioritized: significant percentage of stale fingerprint removal operations (%d%%).

## Parameters

**vvol\_name** (STRING): Volume name.

**high\_verify\_pct** (INT): Percentage of cumulative stale fingerprint removal operations after the fifth iteration of automatic deduplication runs.

## sis.wafiron events

### sis.wafiron.required

### Severity

ALERT

### Description

This message occurs when a requested SIS volume 'undo' operation fails to correct on-disk corruption and requires wafiron to be run.

### Corrective Action

Start wafiron (Command Line: 'aggr wafiron start'. See 'aggr help wafiron' for help) to correct on-disk corruptions after completion of current 'sis undo' operation on the volume.

### Syslog Message

SIS 'undo' command detected on-disk corruption while working on unshare request. Run wafiron for volume '%s%s' of aggregate '%s' after completion of current 'sis undo' operation on the volume. See 'aggr help wafiron' for options on running wafiron.

## Parameters

**volume** (STRING): Name of the affected volume.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**aggregate** (STRING): Aggregate name containing the affected volume.

# siw events

## siw.cpubindingfailed events

### siw.cpuBindingFailed

#### Severity

ERROR

#### Description

This message occurs when the system fails to bind a SW-iWARP thread to a CPU during the boot-up. If this happens, SW-iWARP cannot support polling and will be running in a degraded state.

#### Corrective Action

Try a takeover and giveback of both of the nodes. If it still fails, contact NetApp technical support.

#### Syslog Message

SW-iWARP thread could not be bound to a CPU because: "%s".

#### Parameters

**reason** (STRING): Reason for the CPU binding failure

## siw.mpaverrelaxed events

### siw.mpaVerRelaxed

#### Severity

ERROR

#### Description

This message occurs when there is a mismatch in SW-iWARP versions, but due to negotiation or enhanced connect being enabled. The SW-iWARP booted up with reduced functionality and performance capabilities.

#### Corrective Action

Upgrade the ONTAP® release on the local node and the partner node to the same version of SW-iWARP to run with full functionality and performance capabilities.

#### Syslog Message

SIW Version mismatch between the local node and partner node. Local IRD: %d ORD: %d Partner IRD: %d ORD: %d

#### Parameters

**localIRD** (INT): Local IRD value.

**localORD** (INT): Local ORD value.

**partnerIRD** (INT): Partner IRD value.

**partnerORD** (INT): Partner ORD value.

# siw.mpaversionmismatch events

## siw.mpaVersionMismatch

### Severity

ALERT

### Description

This message occurs when there is SW-iWARP Version mismatch between the HA nodes. SW-iWARP cannot start up and nodes are not connected.

### Corrective Action

Upgrade the ONTAP® release on the nodes to the same version of SW-iWARP to bring up the SW-iWARP and connect the nodes.

### Syslog Message

SW-iWARP version mismatch between the local node and the partner node.

### Parameters

(None).

# siw.pollingenablefailed events

## siw.pollingEnableFailed

### Severity

ERROR

### Description

This message occurs when SW-iWARP device polling cannot be enabled during the boot-up and SW-iWARP will be running in a degraded state.

### Corrective Action

Try a takeover and giveback of the local node. If it still fails, contact NetApp technical support.

### Syslog Message

SW-iWARP polling could not be enabled because the polling device was not ready.

### Parameters

(None).

# sk events

## sk.hog events

### sk.hog.runtime

**Severity**

NOTICE

**Description**

This event is issued when a process runs too long.

**Corrective Action**

(None).

**Syslog Message**

Process %s ran for %d milliseconds

**Parameters**

**procname** (STRING): The name of the process which ran too long

**runtime** (INT): Number of milliseconds process ran

## sk.kt events

### sk.kt.rtc

**Severity**

INFORMATIONAL

**Description**

This is the output of debugging command that prints the time of day held in the real-time clock chip.

**Corrective Action**

(None).

**Syslog Message**

RTC time: %s

**Parameters**

**rtc\_time** (STRING): The time held in the real-time clock chip.

## sk.panic events

### sk.panic

**Severity**

ALERT

## Description

This event is issued when a panic occurs.

## Corrective Action

(None).

## Syslog Message

Panic String: %s

## Parameters

**reason** (STRING): (None).

# sk.ring events

## sk.ring.buffer.saved

### Severity

INFORMATIONAL

### Description

This message indicates that a module detected a condition that warranted dumping the internal tracing buffer to a file.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**module** (STRING): Module requesting the dump.

**filename** (STRING): Filename-extension of the dump file created. The filename-base is accessible by running the 'sysctl ntap.sktrace.filename' command at a shell prompt.

**description** (STRING): Description of the state that caused the dump and when the dumped file should be examined. For example, if the buffers are dumped during replication and the system is experiencing replication errors, then the dumped file might hold information important for debugging.

# sla events

## sla.shelf events

### sla.shelf.mod.reboot

#### Severity

NOTICE

#### Description

This message occurs when a shelf firmware module is rebooted administratively, or when a reboot occurs normally when the system powers up, the firmware is updated, or the alternate control path (ACP) is reset.

#### Corrective Action

(None).

#### Syslog Message

Reboot event reported by module %s in shelf: %s, log: %s

#### Parameters

**module** (STRING): Module designator (A or B).

**shelf** (STRING): Shelf identifier.

**log** (STRING): Reboot event log from the shelf module.

### sla.shelf.mod.reboot.unexp

#### Severity

ERROR

#### Description

This message occurs when a shelf firmware module reboots unexpectedly (not due to powering up, a firmware update, an alternate control path (ACP) reset, nor an administrative reboot).

#### Corrective Action

Check the cabling to the shelf. Check the status LED on the shelf, making sure that the fault LEDs are not on. Clear any fault condition, if possible.

#### Syslog Message

Unexpected reboot event reported by module %s in shelf: %s, log: %s

#### Parameters

**module** (STRING): Module designator (A or B).

**shelf** (STRING): Shelf identifier.

**log** (STRING): Reboot event log from the shelf module.

# sm events

## sm.break events

### sm.break.enable.se

#### Deprecated

Deprecated as of version 9.14.1.

#### Severity

NOTICE

#### Description

This message occurs when the "storage efficiency" feature is enabled on a destination volume due to a broken SnapMirror® relationship on an All Flash Optimized system. When the relationship breaks, the DP volume is converted to RW, and all RW volumes have the "storage efficiency" feature enabled.

#### Corrective Action

(None).

#### Syslog Message

"Storage efficiency" feature has been enabled on the destination volume %s due to a SnapMirror break operation.

#### Parameters

**dst\_path** (STRING): Destination FlexVol® volume name on which the "storage efficiency" feature is enabled.

## sm.c2c events

### sm.c2c.rebaseline.xfer

#### Severity

INFORMATIONAL

#### Description

This message occurs when SnapMirror® performs a SnapMirror Cloud rebaseline transfer because certain conditions for an incremental transfer haven't been met between the source volume and the object-store destination endpoint of the transfer.

#### Corrective Action

(None).

#### Syslog Message

This will be a rebaseline transfer because certain conditions for an incremental transfer haven't been met between the source volume '%s' and the destination object-store endpoint '%s', with endpoint uuid '%s', and relationship uuid '%s'. Refer to the sktrace logs for additional details on which conditions haven't been met.



## Parameters

**src\_path** (STRING): Path name of the source volume of the SnapMirror relationship.

**dst\_ep\_path** (STRING): Path name of the destination object-store endpoint of the SnapMirror relationship.

**dst\_ep\_uuid** (STRING): UUID of the destination object-store endpoint of the SnapMirror relationship.

**rel\_uuid** (STRING): UUID of the SnapMirror relationship.

## sm.c2c.xfer.fail.no.com.snap

### Severity

ERROR

### Description

This message occurs when a SnapMirror® Cloud transfer fails because there is no common Snapshot(tm) copy between the source volume and the object-store destination endpoint of the transfer.

### Corrective Action

Retry the transfer by passing a specific Snapshot copy name using the "-s" option. Note that once you choose to go ahead with that option, you will not be able to transfer Snapshot copies older than the specified Snapshot copy. So it is suggested to select a Snapshot copy that is older than the Snapshot copies that you would like to be backed up per schedule.

### Syslog Message

There is no common Snapshot copy between the source volume '%s' and the destination object-store endpoint '%s', with endpoint uuid '%s', and relationship uuid '%s'.

## Parameters

**src\_path** (STRING): Path name of the source volume of the SnapMirror relationship.

**dst\_ep\_path** (STRING): Path name of the destination object-store endpoint of the SnapMirror relationship.

**dst\_ep\_uuid** (STRING): UUID of the destination object-store endpoint of the SnapMirror relationship.

**rel\_uuid** (STRING): UUID of the SnapMirror relationship.

## sm.cascade events

### SM.cascade.concurrency.unsupported

### Severity

NOTICE

### Description

This message occurs when concurrent SnapMirror® transfers are attempted while in a cascaded scenario, which is not supported. For example, in a relationship A→B→C, when an A→B SnapMirror update is occurring, a B→C SnapMirror update/initialize or a B→C SMTape backup cannot happen at the same time..

### Corrective Action

(None).

### Syslog Message

%s cascade concurrency is not supported for volume %s until the current transfer has ended.

## Parameters

**product** (STRING): Either VSM-SMTape cascade or VSM-VSM/SMTape cascade

**dstpath** (STRING): Replication destination system.

# sm.disable events

## sm.disable.se

### Severity

NOTICE

### Description

This message occurs when the "storage efficiency" feature on a destination volume is automatically disabled successfully by the SnapMirror® resync or update operation.

### Corrective Action

(None).

### Syslog Message

Disabled the "storage efficiency" feature for volume %s during a %s.

### Parameters

**dst\_path** (STRING): Destination FlexVol® volume path on which the "storage efficiency" feature is disabled.

**workflow\_type** (STRING): Workflow that disabled the "storage efficiency feature".

## sm.disable.se.fail

### Severity

ERROR

### Description

This message occurs when the "storage efficiency" feature fails to be disabled on a destination volume during a SnapMirror® resync or update operation.

### Corrective Action

Disable storage efficiency manually by using the "volume efficiency off "volume\_name" " command. If the failure persists, contact NetApp technical support.

### Syslog Message

Failed to disable the "storage efficiency" feature for volume %s during a %s. Reason: %s.

### Parameters

**dst\_path** (STRING): Destination FlexVol® volume path on which the "storage efficiency" feature failed to be disabled.

**workflow\_type** (STRING): Workflow that failed to disable the "storage efficiency" feature.

**reason** (STRING): Reason for the failure.

# sm.dst events

## sm.dst.vol.no.tag.metadata

### Severity

ALERT

### Description

This message occurs when the SnapMirror® labels are lost for all Snapshot(tm) copies on the vault destination. This happens when the Snapshot copy tag's metafile of a volume has an inconsistent header. As a result, SnapMirror cannot automatically delete existing Snapshot copies from the SnapVault® destination and can reach the maximum supported Snapshot copies limit. New Snapshot copies are not affected.

### Corrective Action

When SnapVault destination reaches maximum supported Snapshot copies limit, delete the Snapshot copies that are missing the SnapMirror label. Use "volume snapshot show -snapmirror-label -" to get a list of Snapshot copies that are missing SnapMirror label.

### Syslog Message

The volume '%s' has lost SnapMirror label for all Snapshot copies. SnapMirror will use Snapshot copy's Name Prefix to automatically delete Snapshot copies based on SnapMirror policy '%s'.

### Parameters

**dst\_path** (STRING): Destination volume of the SnapVault relationship.  
**policy\_name** (STRING): SnapMirror policy name.

# sm.fgmirror events

## sm.fgmirror.stale.grp.wkflw

### Severity

ERROR

### Description

This message occurs when improper group workflow cleanup caused the subsequent SnapMirror® operation to fail.

### Corrective Action

Run the "snapmirror abort -h {-relationship-id <relationshipID>}" command, and then retry the SnapMirror operation.

### Syslog Message

SnapMirror operation on relationship with relationshipID %s failed because of an existing stale group workflow.

### Parameters

**relationshipID** (STRING): UUID of the SnapMirror relationship.

# sm.keepcnt events

## sm.keepcnt.reached

### Severity

NOTICE

### Description

This message occurs when new Snapshot(tm) copies on the source volume cannot be transferred because the Snapshot copy retention count (keep count) for Snapshot copies with a given snapmirror-label has been reached on the destination volume and the 'preserve' flag is set to 'true'.

### Corrective Action

See "Working with FlexVol volume SnapVault backups" in the Clustered Data ONTAP Data Protection Guide for details about managing SnapVault relationships.

### Syslog Message

No new Snapshot copies with snapmirror-label(s) %s can be transferred to destination volume %s because the keep count has been reached and the preserve value is set to 'true'.

### Parameters

**snap\_label** (STRING): Comma-separated list of snapmirror-labels for the Snapshot copies that cannot be transferred.

**dst\_path** (STRING): Destination volume of the SnapMirror® relationship.

# sm.log events

## sm.log.write.fail

### Severity

ERROR

### Description

This message occurs when writing SnapMirror® data to a log file fails. The most likely cause is insufficient space on the root volume.

### Corrective Action

If the failure is due to lack of space, increase the root volume size or identify any SnapMirror logs that might not be needed, verify that backups exist for them, and then remove them. To check for available space across the cluster, use the 'df' command. If the condition persists, contact NetApp technical support.

### Syslog Message

Writing log data of size %d bytes to SnapMirror log file '%s' failed with error '%s'.

### Parameters

**size** (INT): Data size in bytes.

**filepath** (STRING): Log file path.

**error** (STRING): Error string.

# sm.mediator events

## sm.mediator.added

### Severity

NOTICE

### Description

This message occurs when ONTAP® Mediator is added successfully on a cluster.

### Corrective Action

(None).

### Syslog Message

ONTAP Mediator (version %s) is added on cluster '%s' having peer cluster '%s' and mediator IP address '%s'.

### Parameters

**version** (STRING): ONTAP Mediator version.

**cluster** (STRING): Cluster name on which ONTAP Mediator was added.

**peerCluster** (STRING): Peer cluster name.

**ipAddress** (STRING): ONTAP Mediator IP address.

## sm.mediator.cacert.expired

### Severity

EMERGENCY

### Description

This message occurs when the ONTAP® Mediator certificate authority (CA) certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.

### Corrective Action

Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

### Syslog Message

CA certificate of the ONTAP Mediator (IP: %s) expired on %s.

### Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.

**expiryDate** (STRING): Expiry date of the ONTAP Mediator CA certificate.

## sm.mediator.cacert.expiring

### Severity

ALERT

## Description

This message occurs when the ONTAP® Mediator certificate authority (CA) certificate is due to expire within the next 30 days.

## Corrective Action

Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

## Syslog Message

CA certificate for the ONTAP Mediator (IP: %s) will expire in %d days. Expiry: %s.

## Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.  
**daysToExpire** (INT): Number of days remaining until the certificate expires.  
**expiryDate** (STRING): Expiry date of the ONTAP Mediator CA certificate.

## sm.mediator.clientc.expired

### Severity

EMERGENCY

### Description

This message occurs when the ONTAP® Mediator client certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.

### Corrective Action

Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

### Syslog Message

Client certificate of the ONTAP Mediator (IP: %s) expired on %s.

### Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.  
**expiryDate** (STRING): Expiry date of the ONTAP Mediator client certificate.

## sm.mediator.clientc.expiring

### Severity

ALERT

### Description

This message occurs when the ONTAP® Mediator client certificate is due to expire within the next 30 days.

### Corrective Action

Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

## Syslog Message

Client certificate for the ONTAP Mediator (IP: %s) will expire in %d days. Expiry: %s.

## Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.  
**daysToExpire** (INT): Number of days remaining until the certificate expires.  
**expiryDate** (STRING): Expiry date of the ONTAP Mediator client certificate.

## sm.mediator.in.quorum

### Severity

NOTICE

### Description

This message occurs when the ONTAP® Mediator connectivity is reestablished, its records are resynchronized and it is ready to participate in SnapMirror® Business Continuity (SMBC) quorum activities.

### Corrective Action

(None).

## Syslog Message

The ONTAP Mediator (IP: %s) connectivity is reestablished, its records are resynchronized and it is ready to participate in SMBC quorum activities. Peer cluster: %s.

## Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.  
**peerCluster** (STRING): Peer cluster name.

## sm.mediator.misconfigured

### Severity

ALERT

### Description

This message occurs when either the ONTAP® Mediator is repurposed or the Mediator package is no longer installed on the Mediator server. As a result, SnapMirror® failover is not possible.

### Corrective Action

Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

## Syslog Message

ONTAP Mediator is not accessible on cluster "%s" with Mediator IP address "%s".

## Parameters

**cluster** (STRING): Name of the cluster that is attempting to access the ONTAP Mediator.  
**ipAddress** (STRING): IP address of the ONTAP Mediator.

## sm.mediator.removed

**Severity**

NOTICE

**Description**

This message occurs when ONTAP® Mediator is removed successfully from a cluster.

**Corrective Action**

(None).

**Syslog Message**

ONTAP Mediator (version %s) was removed on cluster '%s' having peer cluster '%s' and mediator IP address '%s'.

**Parameters**

**version** (STRING): ONTAP Mediator version.

**cluster** (STRING): Cluster name on which ONTAP Mediator was removed.

**peerCluster** (STRING): Peer cluster name.

**ipAddress** (STRING): ONTAP Mediator IP address.

**sm.mediator.removed.force****Severity**

ERROR

**Description**

This message occurs when ONTAP® Mediator is removed from a cluster using the "force" option. As a result, the ONTAP Mediator configuration might still be present on the peer cluster. This results in a partial configuration that could lead to the failure of a SnapMirror® failover operation.

**Corrective Action**

Check whether the current ONTAP Mediator is configured on the peer cluster. If so, remove the ONTAP Mediator by using the (diagnostic privilege) "snapmirror mediator remove -force true" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

**Syslog Message**

ONTAP Mediator (IP: %s, version: %s) was removed using the "force" option on cluster %s. Peer cluster %s.

**Parameters**

**ipAddress** (STRING): ONTAP Mediator IP address.

**version** (STRING): ONTAP Mediator version.

**cluster** (STRING): Cluster name on which ONTAP Mediator was removed.

**peerCluster** (STRING): Peer cluster name.

**sm.mediator.resumed****Severity**

NOTICE

**Description**

This message occurs when ONTAP® Mediator is resumed successfully on a cluster.



**Corrective Action**

(None).

**Syslog Message**

ONTAP Mediator was resumed on cluster '%s' having mediator IP address '%s'.

**Parameters**

**cluster** (STRING): Cluster name on which ONTAP Mediator was resumed.

**ipAddress** (STRING): ONTAP Mediator IP address.

**sm.mediator.serverc.expired****Severity**

EMERGENCY

**Description**

This message occurs when the ONTAP® Mediator server certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.

**Corrective Action**

Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

**Syslog Message**

Server certificate of the ONTAP Mediator (IP: %s) expired on %s.

**Parameters**

**ipAddress** (STRING): ONTAP Mediator IP address.

**expiryDate** (STRING): Expiry date of the ONTAP Mediator server certificate.

**sm.mediator.serverc.expiring****Severity**

ALERT

**Description**

This message occurs when the ONTAP® Mediator server certificate is due to expire within the next 30 days.

**Corrective Action**

Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

**Syslog Message**

Server certificate for the ONTAP Mediator (IP: %s) will expire in %d days. Expiry: %s.

**Parameters**

**ipAddress** (STRING): ONTAP Mediator IP address.

**daysToExpire** (INT): Number of days remaining until the certificate expires.

**expiryDate** (STRING): Expiry date of the ONTAP Mediator server certificate.

## **sm.mediator.suspended**

### **Severity**

NOTICE

### **Description**

This message occurs when ONTAP® Mediator is suspended successfully on a cluster.

### **Corrective Action**

(None).

### **Syslog Message**

ONTAP Mediator was suspended on cluster '%s' having mediator IP address '%s'.

### **Parameters**

**cluster** (STRING): Cluster name on which ONTAP Mediator was suspended.

**ipAddress** (STRING): ONTAP Mediator IP address.

## **sm.mediator.unreachable**

### **Severity**

ALERT

### **Description**

This message occurs when the ONTAP® Mediator is unreachable on a cluster. As a result, SnapMirror® failover is not possible.

### **Corrective Action**

Check the network connectivity to the ONTAP Mediator by using the "network ping" and "network traceroute" commands. If the issue persists, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

### **Syslog Message**

ONTAP Mediator (IP: %s) is unreachable from cluster %s.

### **Parameters**

**ipAddress** (STRING): IP address of the ONTAP Mediator.

**cluster** (STRING): Name of the cluster that is attempting to access the ONTAP Mediator.

## **sm.mediator.unusable**

### **Severity**

ERROR

### **Description**

This message occurs when ONTAP® Mediator is unusable due to a corrupted mediator server database. If ONTAP Mediator is in this state, SnapMirror® failover will not be possible.

## Corrective Action

Remove ONTAP Mediator by using the "snapmirror mediator remove" command, and then add it again by using the "snapmirror mediator add" command.

## Syslog Message

ONTAP Mediator is unusable on cluster '%s' having mediator IP address '%s'.

## Parameters

**cluster** (STRING): Cluster name on which ONTAP Mediator is unusable.

**ipAddress** (STRING): ONTAP Mediator IP address.

# sm.no events

## sm.no.metadata.objstore

### Severity

ALERT

### Description

This message occurs when a SnapMirror® catalog metadata transfer is attempted and a SnapMirror object store configuration with "usage" of "metadata" is not configured on the system. Catalog metadata transfers will fail until an appropriate object store is configured.

### Corrective Action

Either add "usage" of "metadata" to an existing SnapMirror object store configuration or create a new SnapMirror object store configuration with "usage" of "metadata". Use the "snapmirror object-store config modify -usage" command or the "snapmirror object-store config create" command to modify or create an object store configuration.

### Syslog Message

No SnapMirror object store configuration with "usage" of "metadata" found for catalog metadata transfer to destination '%s'. Catalog metadata transfers will fail until an appropriate object store is configured.

### Parameters

**dst\_path** (STRING): Destination of the SnapMirror relationship.

# sm.slc events

## sm.slc.async.resync.preserve

### Severity

NOTICE

### Description

This message occurs when a SnapMirror® resync operation detects data divergence between the common Snapshot(tm) copy and the AFS on a resync destination volume where the destination volume is a SnapLock Compliance volume and the SnapMirror relationship is of "async-mirror" policy type. The resync operation preserves the data changes in a locked Snapshot copy for the duration of the current volume expiry time. If the volume expiry time is in the past or has not been set, then the Snapshot copy is locked for a duration of 30 days. The common Snapshot copy is also locked for the same duration.

## Corrective Action

(None).

## Syslog Message

SnapMirror resync operation created a new Snapshot copy '%s' to preserve the data changes in the AFS with reference to common Snapshot copy '%s' at destination SnapLock Compliance volume '%s'. Both Snapshot copies are locked for the duration of the current volume expiry time '%s'. If the volume expiry time is in the past or has not been set, then the Snapshot copies are locked for a duration of 30 days.

## Parameters

**new\_snap\_name** (STRING): Name of the new Snapshot copy; where the data changes are preserved.

**common\_snap\_name** (STRING): Name of the common Snapshot copy on the destination volume.

**dst\_path** (STRING): Destination volume of the SnapMirror relationship.

**dst\_vol\_expiry** (STRING): Expiry time of the destination volume.

**relationship\_id** (STRING): UUID of the SnapMirror relationship.

# sm.syncmirror events

## sm.syncmirror.out.of.sync

### Severity

ALERT

### Description

This message occurs when a sync granular CG relationship transitions from in-sync status to out-of-sync status.

### Corrective Action

Verify network connections between the primary cluster and the secondary cluster. Check for error information in the "snapmirror show" command output for the relationship. Quiesce the relationship if the condition cannot be addressed immediately.

### Syslog Message

Sync granular CG relationship with source path %s and destination path %s has transitioned from in-sync to out-of-sync for the follow reason: %s.

### Parameters

**src\_path** (STRING): Source path of the sync granular CG relationship that transitioned from in-sync to out-of-sync.

**dst\_path** (STRING): Destination path of the sync granular CG relationship that transitioned from in-sync to out-of-sync.

**reason** (STRING): Reason that the relationship went out-of-sync.

# sm.vfiler events

## SM.vfiler.config.move.warn

### Severity

NOTICE

## Description

This message occurs when movement of the SnapMirror® status configuration from 'from' vFiler to 'to' vFiler fails. The SnapMirror or SnapVault® status might show the wrong status entries because the status entries are not copied.

## Corrective Action

(None).

## Syslog Message

Movement of SnapMirror status configuration failed while moving volumes between vFiler %s to vFiler %s. SnapMirror or SnapVault status might show the wrong status entries.

## Parameters

**from\_vfiler** (STRING): vFiler(tm) from which resources are being moved.

**to\_vfiler** (STRING): vFiler to which resources are being moved.

# sm.xfer events

## sm.xfer.afs.divergence

### Severity

NOTICE

### Description

This message occurs when active file system diverges on a SnapMirror® destination. The active file system version UUID must match the version UUID of the latest Snapshot(tm) copy on the destination volume or the Snapshot copy that is currently being replicated to the destination volume. The system automatically resolves the divergence by performing a SnapRestore® operation.

### Corrective Action

(None).

### Syslog Message

Active file system with version UUID %s has diverged from the latest Snapshot copy %s version UUID %s on destination volume %s.

### Parameters

**afs\_uuid** (STRING): Version UUID of the active file system. Internal use only.

**snap\_name** (STRING): Name of the latest Snapshot copy on the destination volume.

**latest\_snap\_uuid** (STRING): Version UUID of the latest Snapshot copy on the destination volume. Internal use only.

**dst\_path** (STRING): Destination volume of the SnapMirror relationship.

## sm.xfer.fail.del.snap

### Severity

ERROR

### Description

This message occurs when deletion of a Snapshot(tm) copy on the destination volume fails, and therefore a subsequent transfer exceeds the maximum supported Snapshot copy count for the vault destination.

### Corrective Action

Manually delete the Snapshot copies from the destination volume and reissue the failed SnapMirror command.

### Syslog Message

Unable to delete Snapshot copy '%s' from destination volume '%s' of the relationship with relationship UUID '%s' to make space for for newer Snapshot copies.

### Parameters

**snap\_name** (STRING): Snapshot copy that was not deleted.

**dst\_path** (STRING): Destination volume of the relationship.

**relationship\_id** (STRING): UUID of the snapmirror relationship.

## sm.xfer.fail.maxlim.snap

### Severity

ALERT

### Description

This message occurs when a SnapMirror® transfer fails because the maximum supported Snapshot(tm) copy limit for a SnapMirror destination is reached and there is no space for a new, incoming Snapshot copy.

### Corrective Action

Reduce the "keep" count by using the "snapmirror policy modify-rule -policy [policy\_name] -keep [keep\_count]" command, or by manually deleting Snapshot copies from the SnapMirror® destination.

### Syslog Message

The number of Snapshot copies on destination volume '%s' of the relationship with relationship UUID '%s' and policy '%s' has reached the maximum supported.

### Parameters

**dst\_path** (STRING): Destination volume of the relationship.

**relationship\_id** (STRING): UUID of the snapmirror relationship.

**policy\_name** (STRING): SnapMirror policy name.

## sm.xfer.incr.snap.rename

### Severity

ERROR

### Description

This message occurs when an incomplete SnapMirror® operation is detected for an incremental Snapshot(tm) copy that no longer exists on the SnapMirror® source volume. As a result, SnapMirror creates a Snapshot copy with the name "snapmirror\_vlt\_xfer\_recovery" and requires a SnapMirror Resync operation to be performed on the SnapVault relationship to resume data protection.

### Corrective Action

Resume data protection either by retaining or deleting the Snapshot(tm) copy. Use the 'snapmirror resync -preserve' to retain it or use the 'snapmirror resync' command to delete it.

## Syslog Message

An incomplete SnapMirror operation was detected for an incremental Snapshot copy that no longer exists on the source volume '%s'. The data transfer for the incremental Snapshot copy was completed but the Snapshot copy on the vault destination '%s' was not created. Instead, a Snapshot copy with the name "snapmirror\_vlt\_xfer\_recovery" was created.

## Parameters

**src\_path** (STRING): Source volume of the relationship.

**dst\_path** (STRING): Destination volume of the relationship.

## sm.xfer.no.new.snap

### Severity

INFORMATIONAL

### Description

This message occurs when a SnapMirror® update does not find any new Snapshot(tm) copies to transfer. This can happen when none of the newer Snapshot copies on the source has a SnapMirror® label that matches a SnapMirror policy.

### Corrective Action

Check that the SnapMirror label from the Snapshot policy matches SnapMirror policy rules.

## Syslog Message

SnapMirror update did not find any new Snapshot copies to transfer from the source volume '%s' of the relationship with relationship UUID '%s'.

## Parameters

**src\_path** (STRING): Source volume of the SnapMirror relationship.

**relationship\_id** (STRING): UUID of the snapmirror relationship.

## sm.xfer.warn.limit

### Severity

ERROR

### Description

This message occurs when a SnapMirror® transfer receives a warning when it is approaching the "keep" count of a particular SnapMirror® relationship.

### Corrective Action

Increase the "keep" count to transfer more Snapshot(tm) copies with the same label using the following command: snapmirror policy modify-rule -policy [policy\_name] -keep [keep\_count]

## Syslog Message

Warnings issued for SnapMirror relationship with source volume '%s' and destination volume '%s' with relationship UUID '%s' : %s

## Parameters

**src\_path** (STRING): Source volume of the relationship.

**dst\_path** (STRING): Destination volume of the relationship.

**relationship\_id** (STRING): UUID of the snapmirror relationship.

**warn\_str** (STRING): SnapMirror transfer warning.

## **sm.xfer.warn.snapshot.count.limit**

### **Severity**

ERROR

### **Description**

This message occurs during a SnapMirror® transfer, when the Snapshot(tm) count limit is reached for the SnapMirror relationship between the indicated source and destination volumes.

### **Corrective Action**

Delete some Snapshot copies from the destination volume to create space for new Snapshot copies to be transferred to it.

### **Syslog Message**

Snapshot count limit for SnapMirror relationship with source volume '%s' and destination volume '%s' with relationship UUID '%s' has been reached : %s

### **Parameters**

**src\_path** (STRING): Source volume of the SnapMirror relationship.

**dst\_path** (STRING): Destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): UUID of the snapmirror relationship.

**warn\_str** (STRING): SnapMirror transfer warning.



# smbc events

## smbc.aufo events

### smbc.aufo.completed

**Severity**

NOTICE

**Description**

This message occurs when the SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation completes.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror automatic failover completed. Destination path: %s.

**Parameters**

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

### smbc.aufo.during.expand

**Severity**

NOTICE

**Description**

This message occurs when a SnapMirror® Business Continuity (SM-BC) automatic unplanned failover operation starts while a SnapMirror Business Continuity expand operation is in progress. The new volumes of the consistency group that are going through the expand operation do not participate in the failover and remain as data protection (DP) after completion of failover.

**Corrective Action**

(None).

**Syslog Message**

Automatic unplanned failover has started on the SnapMirror Business Continuity relationship while the relationship is expanding. The new volumes of the consistency group that are going through the expand operation do not participate in the failover and remain as DP after completion of failover.

**Parameters**

**dstpath** (STRING): Path name for the destination consistency group of SM-BC relationship.

### smbc.aufo.failed

**Severity**

ERROR

## Description

This message occurs when the SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation fails.

## Corrective Action

Perform "snapmirror delete -destination-path destination\_path" followed by "snapmirror break" on the volumes.

## Syslog Message

SnapMirror automatic failover failed. Destination path: %s.

## Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## smbc.aufo.failover.incapable

### Severity

ERROR

### Description

This message occurs when the SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation fails because the destination site is not failover capable.

### Corrective Action

Before applying corrective action, the User must ensure that the primary/master copy is unavailable, otherwise the assumption is that Master copy is actively being used by the Application workload. Perform "snapmirror delete -destination-path destination\_path" followed by "snapmirror break" on the volumes.

### Syslog Message

SnapMirror automatic failover encountered a failure because the site with destination path "%s" is not failover-capable.

### Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## smbc.aufo.io.outage.complete

### Severity

NOTICE

### Description

This message occurs when I/O can resume after an automatic unplanned failover of a SnapMirror Business Continuity (SMBC) relationship.

### Corrective Action

(None).

### Syslog Message

In the automatic unplanned failover of the SMBC relationship with destination path %s and relationship-ID %s, I/O disruption started at %s and lasted for %u msecs. Mediator Agent processing started at %s and lasted for %u msecs.

## Parameters

**dstpath** (STRING): Destination path of the SMBC relationship.

**cg\_rel\_id** (STRING): ID of the SMBC relationship.

**io\_outage\_start\_time** (STRING): I/O disruption start time.

**io\_outage\_latency** (LONGINT): I/O disruption window (msecs).

**tba\_trigger\_start\_time** (STRING): Mediator Agent processing start time.

**tba\_trigger\_latency** (LONGINT): Mediator Agent processing window (msecs).

## **smbc.aufo.queued**

### Severity

NOTICE

### Description

This message occurs when the SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation is queued.

### Corrective Action

(None).

### Syslog Message

SnapMirror Business Continuity (SMBC) automatic unplanned operation failover queued. Destination path: %s.

### Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## **smbc.aufo.start**

### Severity

NOTICE

### Description

This message occurs when the SnapMirror® Business Continuity (SMBC) automatic unplanned failover operation starts.

### Corrective Action

(None).

### Syslog Message

SnapMirror automatic failover started. Destination path: %s.

### Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## **smbc.avail events**

### **smbc.avail.aufoDisabled**

**Severity**

ERROR

**Description**

This message occurs when the primary-bias state is set for the secondary endpoints of the SnapMirror Business Continuity (SMBC) relationships because both of the SMBC clusters lose connectivity to the ONTAP Mediator. It prevents an automatic unplanned failover operation on the secondary endpoints.

**Corrective Action**

Check the network connection and restore the access to the ONTAP Mediator.

**Syslog Message**

Primary-bias state is set for the secondary endpoints of the SMBC relationships. ONTAP Mediator IP: "%s", Peer cluster: "%s".

**Parameters**

**ipAddress** (STRING): ONTAP Mediator IP address.

**peerCluster** (STRING): Peer cluster name.

**smbc.avail.aufoEnabled****Severity**

NOTICE

**Description**

This message occurs when the primary-bias state is cleared for the secondary endpoints of the SnapMirror Business Continuity (SMBC) relationships. It resumes the ONTAP Mediator assisted automatic unplanned failover operation on secondary endpoints.

**Corrective Action**

(None).

**Syslog Message**

Primary-bias state is cleared for the secondary endpoints of the SMBC relationships. ONTAP Mediator IP: "%s", Peer cluster: "%s".

**Parameters**

**ipAddress** (STRING): ONTAP Mediator IP address.

**peerCluster** (STRING): Peer cluster name.

**smbc.avail.primaryBias****Severity**

ERROR

**Description**

This message occurs when the primary-bias state is set for the primary endpoints of the SnapMirror Business Continuity (SMBC) relationships because both of the SMBC clusters lose connectivity to the ONTAP Mediator. It allows primary endpoints to drive the consensus locally without depending on the ONTAP Mediator.

### Corrective Action

Check the network connection and restore the access to the ONTAP Mediator.

### Syslog Message

Primary-bias state is set for the primary endpoints of the SMBC relationships. ONTAP Mediator IP: "%s", Peer cluster: "%s".

### Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.

**peerCluster** (STRING): Peer cluster name.

## smbc.avail.primaryBiasReset

### Severity

NOTICE

### Description

This message occurs when the primary-bias state is cleared for the primary endpoints of the SnapMirror Business Continuity (SMBC) relationships. It resumes the three-way quorum involving the ONTAP Mediator and peer cluster to derive consensus.

### Corrective Action

(None).

### Syslog Message

Primary-bias state is cleared for the primary endpoints of the SMBC relationships. ONTAP Mediator IP: "%s", Peer cluster: "%s".

### Parameters

**ipAddress** (STRING): ONTAP Mediator IP address.

**peerCluster** (STRING): Peer cluster name.

## smbc.cg events

### smbc.cg.info.read.failed

### Severity

ERROR

### Description

This message occurs when read of Consistency Group aggregate information entry fails for the SnapMirror® Business Continuity (SMBC) volume which results into IO disruption.

### Corrective Action

To recover from this situation, perform below steps : Run wafiron on the aggregate hosting the SMBC volume which failed during mount time. Run resync on all the SMBC CG relationships present in the aggregate on which wafiron executed.

### Syslog Message

Read of Consistency Group aggregate information entry failed for source volume DSID: "%llu" during "%s" because Consistency Group aggregate information file is corrupted.

## Parameters

**source\_volume\_dsid** (LONGINT): Source Volume DSID

**operation** (STRING): Name of operation

## smbc.cg.info.update.failed

### Severity

ERROR

### Description

This message occurs when add or delete of Consistency Group aggregate information entry fails for the SnapMirror® Business Continuity (SMBC) volume which results into IO disruption.

### Corrective Action

To recover from this situation, perform below steps : Run wafiron on the aggregate hosting the SMBC volume which failed during mount time. Run resync on all the SMBC CG relationships present in the aggregate on which wafiron executed.

### Syslog Message

"%s" of Consistency Group aggregate information entry failed for source volume DSID: "%llu" because Consistency Group aggregate information file is corrupted.

## Parameters

**metafile\_update\_type** (STRING): Name of aggregate metafile update type

**volume\_dsid** (LONGINT): Volume DSID

## smbc.pfo events

### smbc.pfo.completed

#### Severity

NOTICE

#### Description

This message occurs when the SnapMirror® Business Continuity (SMBC) planned failover operation completes.

#### Corrective Action

(None).

#### Syslog Message

SMBC planned failover operation completed. Destination path: %s.

#### Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

### smbc.pfo.failed

#### Severity

ERROR

## Description

This message occurs when the SnapMirror® Business Continuity (SMBC) planned failover operation fails.

## Corrective Action

Determine the cause of the failure by using the "snapmirror failover show -fields error-reason" command. After addressing the cause, retry the "snapmirror failover start -destination-path destination\_path" command.

## Syslog Message

SMBC planned failover operation failed. Destination path: %s.

## Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## smbc.pfo.io.outage.compl.dst

## Severity

NOTICE

## Description

This message occurs when the secondary (erstwhile) side conversion completes after a planned failover operation of a SnapMirror Business Continuity (SMBC) relationship.

## Corrective Action

(None).

## Syslog Message

In the planned failover of the SMBC relationship with destination path %s and relationship-ID %s, the destination side I/O disruption started at %s and lasted for %u msecs.

## Parameters

**dstpath** (STRING): Destination path of the SMBC relationship.

**cg\_rel\_id** (STRING): ID of the SMBC relationship.

**start\_time** (STRING): I/O disruption start time.

**latency** (LONGINT): I/O disruption window (msecs).

## smbc.pfo.io.outage.compl.src

## Severity

NOTICE

## Description

This message occurs when the primary (erstwhile) side conversion completes after a planned failover operation of a SnapMirror Business Continuity (SMBC) relationship and IO outage completes as the Report Target Port Group(RTPG) is consistent.

## Corrective Action

(None).

## Syslog Message

In the planned failover of the SMBC relationship with destination path %s and relationship-ID %s, the

source side I/O disruption started at %s and lasted for %u msecs.

### Parameters

**dstpath** (STRING): Destination path of the SMBC relationship.  
**cg\_rel\_id** (STRING): ID of the SMBC relationship.  
**start\_time** (STRING): I/O disruption start time.  
**latency** (LONGINT): I/O disruption window (msecs).

## smbc.pfo.start

### Severity

NOTICE

### Description

This message occurs when the SnapMirror® Business Continuity (SMBC) planned failover operation starts.

### Corrective Action

(None).

### Syslog Message

SMBC planned failover operation started. Destination path: %s.

### Parameters

**dstpath** (STRING): Path name of the destination group of the SMBC relationship.

## smbc.policy events

### smbc.policy.reassign

### Severity

NOTICE

### Description

This message occurs when the system detects during a flip on a SnapMirror® Sync relationship that the custom policy has been deleted from the source cluster and automatically falls back to the system policy of type "automated-failover".

### Corrective Action

(None).

### Syslog Message

The custom policy has been deleted from the source cluster so the SnapMirror Sync relationship with destination path: %s has been assigned "AutomatedFailover" policy.

### Parameters

**dstpath** (STRING): Path name of the destination group of the SnapMirror Business Continuity relationship.



# smc events

## smc.cg events

### smc.cg.snap.ren.revert.fail

#### Severity

ERROR

#### Description

This message occurs when a FlexGroup® Snapshot copy rename fails and the revert to the original Snapshot copy name also fails.

#### Corrective Action

Manually rename the Snapshot copy by issuing the "snapshot rename" command.

#### Syslog Message

Unable to revert Snapshot copy name from the new name, "%s" to oldname "%s" in volume "%s".

#### Parameters

**newsnapname** (STRING): New Snapshot copy name.

**origsnapname** (STRING): Original Snapshot copy name.

**dstpath** (STRING): Path name of the destination volume to which the Snapshot copy rename was issued.

**error** (INT): Error code of the failure.

## smc.dst events

### smc.dst.snapshot.op

#### Severity

NOTICE

#### Description

This message occurs when SnapMirror creates or deletes a Snapshot copy with a SnapMirror label on a FlexVol destination volume of a SnapMirror® relationship.

#### Corrective Action

(None).

#### Syslog Message

SnapMirror Snapshot %s %s on destination volume %s, label %s, instance UUID %s, version UUID %s, relationship UUID %s.

#### Parameters

**operation** (STRING): Snapshot copy operation either "create" or "delete".

**snapname** (STRING): Snapshot copy name.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**snaplabel** (STRING): SnapMirror label of the Snapshot copy.

**instance\_uuid** (STRING): Instance UUID of the Snapshot copy.

**version\_uuid** (STRING): Version UUID of the Snapshot copy.

**relationshipID** (STRING): UUID of the SnapMirror relationship.

## smc.frest events

### smc.frest.snap.delete.fail

#### Severity

ERROR

#### Description

This message occurs when a SnapMirror file restore operation fails to delete a local file restore Snapshot copy on the destination volume.

#### Corrective Action

Manually delete snapshot by issuing snapshot delete command.

#### Syslog Message

Unable to delete Snapshot copy '%s' from volume '%s'.

#### Parameters

**snapname** (STRING): Snapshot copy that was not deleted.

**dstpath** (STRING): Path name of the destination volume to which restore was issued.

## smc.restore events

### smc.restore.src.cleanup.fail

#### Severity

ERROR

#### Description

This message occurs when a SnapMirror restore operation or a SnapMirror restore cleanup operation fails to cleanup a SnapMirror restore source relationship information. This issue can happen regardless of the success or failure of the SnapMirror operation.

#### Corrective Action

If the SnapMirror restore operation or the SnapMirror restore cleanup operation has failed, check the reason for the error, take action accordingly, and then run the command again. If the operation has succeeded, run the "snapmirror list-destinations -type RST" command on the cluster containing the source information for the SnapMirror restore relationship, and identify the corresponding relationship. Then remove the stale entry by running the "snapmirror release -destination-path <dst\_path> -source-path <src\_path>" command.

#### Syslog Message

Restore of volume '%s' from volume '%s' with relationship UUID '%s' failed to remove source information for the relationship.

#### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror restore relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror restore relationship.

**relationshipID** (STRING): Relationship UUID.

# smc.snap events

## smc.snap.upd.retention.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror operation fails to copy the retention period of a Snapshot copy from the source endpoint to the destination endpoint because the ONTAP® version of the destination cluster does not support copying the retention period. This issue can happen regardless of the success or failure of the SnapMirror operation.

### Corrective Action

Check the ONTAP version of the destination cluster using the command "system image show -node <node\_name> -iscurrent true -fields version". If the ONTAP version of the destination cluster is earlier than ONTAP 9.12.1 then copying the retention period is not supported on the destination cluster. The destination cluster should be upgraded to ONTAP 9.12.1 or later to copy the retention period of Snapshot copies from the source to the destination.

### Syslog Message

The SnapMirror relationship with relationship UUID "%s" will not copy the retention period of a Snapshot copy from the source volume "%s" to the destination volume "%s" because the ONTAP version of the destination cluster does not support copying the retention period.

### Parameters

**relationshipID** (STRING): Relationship UUID.

**srcpath** (STRING): Path name of the source endpoint of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination endpoint of the SnapMirror relationship.

# smc.snapmir events

## smc.snapmir.dst.invalidType

### Severity

ERROR

### Description

This message occurs when the destination volume of a SnapMirror® relationship is not of -type Data Protection (DP) and a 'snapmirror initialize' operation is attempted. A volume must be of -type DP to accept a 'snapmirror initialize' operation.

### Corrective Action

Use a different volume of -type DP, or use the 'volume create' command to create a new volume of -type DP to use as a destination of the SnapMirror relationship. You cannot modify the type of a volume using the 'volume modify' command.

### Syslog Message

Destination volume '%s' of a SnapMirror relationship is not of -type DP. Relationship UUID '%s'.

## Parameters

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationshipID** (STRING): UUID of the snapmirror relationship.

## smc.snapmir.frest.clean.fail

### Severity

ERROR

### Description

This message occurs when SnapMirror® fails to clean up fenced files and file restore restartability information on a volume and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly and issue 'snapmirror restore -file-restore-clean-up true' on the volume in diagnostic mode.

### Syslog Message

File restore clean up on volume '%s' failed with error '%s'.

## Parameters

**dstpath** (STRING): Path name of the volume.

**error** (STRING): Error message.

## smc.snapmir.init.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror® 'initialize' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

### Syslog Message

Initialize from source volume '%s' to destination volume '%s' failed with error '%s'. Relationship UUID '%s'.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the initialize failed.

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.lrb.auto.resync

### Severity

INFORMATIONAL

## Description

This message occurs when the latest Snapshot® copy on the SnapMirror® or SnapVault® destination is not the base Snapshot copy and the source volume is read-only. Generally this is a rare event and a local-rollback-based SnapMirror resync will be attempted to automatically recover from this situation. If this event occurs frequently, it is likely that the Snapshot policy on the primary volume in the cascade is deleting Snapshot copies too soon.

## Corrective Action

Check the Snapshot policy on the primary (read-write) volume and increase the retention count of the Snapshot copies that must be backed up all the way to the destination.

## Syslog Message

Latest Snapshot copy %s on the destination %s is not the base Snapshot copy. Attempting a local rollback based SnapMirror resync to correct this issue.

## Parameters

**snap\_name** (STRING): Name of the latest Snapshot copy on the SnapMirror or SnapVault destination.  
**dst\_path** (STRING): Destination volume of the SnapMirror/SnapVault relationship.

## smc.snapmir.partnr.noconnect

### Severity

ERROR

### Description

This message occurs when a SnapMirror® Sync relationship cannot detect the heartbeat signal from the partner node.

### Corrective Action

Check the network connection between the source and destination volumes by using the "net interface show -role intercluster -status-admin up" command.

### Syslog Message

Source volume "%s" and destination volume "%s" with relationship UUID "%s" lost heartbeat to partner node.

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.  
**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.  
**relationship\_id** (STRING): Relationship UUID.

## smc.snapmir.release.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror® 'release' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

## Syslog Message

Release on the source volume '%s' for destination volume '%s' failed with error '%s'. Relationship UUID '%s'.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the resync failed.

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.restore.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror® 'restore' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

## Syslog Message

Restore of volume '%s' from volume '%s' failed with error '%s'. Relationship UUID '%s'.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the resync failed.

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.resync.fail

### Severity

ERROR

### Description

This message occurs when a SnapMirror® 'resync' command fails and no more retries will be attempted.

### Corrective Action

Check the reason for the error, take action accordingly, and issue the command again.

## Syslog Message

Resync from source volume '%s' to destination volume '%s' failed with error '%s'. Relationship UUID '%s'.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the resync failed.

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.schd.trans.overrun

### Severity

ERROR

### Description

This message occurs when a scheduled SnapMirror® transfer overruns the schedule window.

### Corrective Action

Check the reason for the schedule overrun and take action accordingly.

### Syslog Message

Scheduled transfer from source volume '%s' to destination volume '%s' is taking longer than the schedule window. Relationship UUID '%s'.

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the scheduled transfer.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.suspending.io

### Severity

ALERT

### Description

This message occurs when SnapMirror® I/O has been suspended, due either to a loss of connectivity between the SnapMirror Mediator software and partner, or to a denial of network access by the Mediator software.

### Corrective Action

Check the status of the connection to the Mediator software by using the "net interface show -role intercluster -status-admin up" command.

### Syslog Message

Suspending I/O as with relationship UUID "%s".

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

**error\_msg** (STRING): Description of the error responsible for the suspension.

## smc.snapmir.tb.noconnect

### Severity

ERROR

### Description

This message occurs when an ONTAP® node is not reachable from the SnapMirror® Mediator software.

## Corrective Action

Check the network connection between the source and destination volumes, by using the "network ping -node node\_name" command.

## Syslog Message

Connection lost to the Mediator software from node: %s, with relationship UUID "%s".

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

**node** (STRING): Node name.

**mediator\_id** (STRING): IP of the Mediator.

## smc.snapmir.unexpected.err

### Severity

ERROR

### Description

This message occurs when a SnapMirror® operation encounters an unexpected error.

### Corrective Action

Retry the SnapMirror operation. If the problem persists, contact NetApp technical support.

### Syslog Message

SnapMirror unexpected error '%s'. Relationship UUID '%s'. Operation UUID '%s'.

### Parameters

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

**operationID** (STRING): Operation UUID.

## smc.snapmir.unsupp.nvmf

### Severity

ERROR

### Description

This message occurs when a SnapMirror® operation encounters an NVMe Namespace that is not supported on the destination side.

### Corrective Action

Upgrade the destination to a version that supports NVMe Namespace objects.

### Syslog Message

SnapMirror unsupported NVMe Namespace. Relationship UUID '%s'.

### Parameters

**relationshipID** (STRING): Relationship UUID.



## smc.snapmir.unsupp.nvmf.ns

### Severity

ERROR

### Description

This message occurs when a SnapMirror® operation encounters a 512-byte block size NVMe namespace that is not supported on the destination.

### Corrective Action

Upgrade the destination to a version that supports 512-byte block size NVMe namespace objects.

### Syslog Message

SnapMirror unsupported 512-byte block size NVMe namespace. Relationship UUID '%s'.

### Parameters

**relationshipID** (STRING): Relationship UUID.

## smc.snapmir.update.fail

### Severity

ERROR

### Description

This message occurs when a scheduled SnapMirror® transfer fails.

### Corrective Action

Check the snapmirror\_error log in the destination filer and take action accordingly.

### Syslog Message

Snapmirror update from source volume '%s' to destination volume '%s' failed with error '%s'. Relationship UUID '%s'.

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the snapmirror transfer failed.

**error** (STRING): Error message.

**relationshipID** (STRING): Relationship UUID.

## smc.snapmirror events

### smc.snapmirror.xfer.complete

### Severity

INFORMATIONAL

### Description

This message occurs when a SnapMirror® transfer completes successfully.

## Corrective Action

(None).

## Syslog Message

Snapmirror transfer from source volume '%s' to destination volume '%s' is successful. '%s'. Relationship UUID '%s'.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume to which the snapmirror transfer failed.

**info** (STRING): Transfer information.

**relationshipID** (STRING): Relationship UUID.

# smc.tot events

## smc.tot.bytes.write.fail

### Severity

NOTICE

### Description

This message occurs when SnapMirror is unable to update the latest total transfer bytes value of the SnapMirror relationship. The total transfer bytes for the relationship might be incorrect.

### Corrective Action

(None).

### Syslog Message

SnapMirror is unable to update the latest total transfer bytes value for the relationship with relationship UUID "%s" and destination "%s".

### Parameters

**relationshipID** (STRING): Relationship UUID.

**dstPath** (STRING): Path name of the destination endpoint of the SnapMirror relationship.

# smc.vlt events

## smc.vlt.cascade.xfer.fail

### Severity

ERROR

### Description

This message occurs when the transfer for the SnapVault® relationship replicating the read-only source volume fails because the SnapMirror® policy prevents it from transferring the exported Snapshot® copy from the source. Generally this happens when the "keep" value of the rule associated with the SnapMirror label of the exported Snapshot copy is reached and the "preserve" flag is set to "true".

**Corrective Action**

Modify the rule for the SnapMirror policy associated with the SnapMirror label of the exported Snapshot copy on the source volume and either set the "preserve" flag to "false" or increase the "keep" count.

**Syslog Message**

Unable to transfer any Snapshot copies to %s because the exported Snapshot copy %s with snapmirror-label %s has a rule whose "keep" value has been reached and the "preserve" flag is set to "true".

**Parameters**

**dst\_path** (STRING): Destination volume of the SnapVault relationship.

**snap\_name** (STRING): Name of the exported Snapshot copy on the source volume.

**snapmirror\_label** (STRING): SnapMirror label associated with the exported Snapshot copy.

# sms events

## sms.cg events

### sms.cg.cascade.not.supported

#### Severity

ERROR

#### Description

This message occurs when a SnapMirror® Synchronous Automated Failover relationship is part of a cascade configuration. Cascade configuration is not supported for SnapMirror® Synchronous Automated Failover relationships.

#### Corrective Action

Delete the Asynchronous cascade leg from the destination of the SnapMirror® Synchronous Consistency Group relationship. Attach the cascade Leg to the source of SnapMirror® Synchronous Consistency Group relationship to re-establish the 3 site DR.

#### Syslog Message

The source of the relationship %s (msid:%llu and vserver\_uuid:%s) is part of of SnapMirror® Synchronous Automated Failover destination relationship. Since cascade is not supported for SnapMirror® Synchronous Automated Failover relationship, failing the operation on the cascade leg.

#### Parameters

**RelationshipId** (STRING): UUID of the SnapMirror Relationship.  
**volumeMSID** (LONGINT): Volume MSID.  
**vserverUUID** (STRING): vserver instance UUID.

## sms.clone events

### sms.clone.request.disallow

#### Severity

NOTICE

#### Description

This message occurs when SnapMirror® Synchronous fails a clone request because the relation is temporarily in a state that does not permit clone replication.

#### Corrective Action

(None).

#### Syslog Message

Clone replication not permitted while in transitioning state.

#### Parameters

(None).

# sms.common events

## sms.common.snapshot.failed

### Severity

ERROR

### Description

This message occurs when there is a failure in creating a common Snapshot(tm) copy. The SnapMirror® Sync relationship continues to be in "in-sync" status. The latest common Snapshot copy is used for recovery in case the relationship status changes to "out-of-sync." The common Snapshot copy should be created at scheduled intervals to decrease the recovery time of "out-of-sync" relationships.

### Corrective Action

Create a common Snapshot copy manually by using the "snapmirror update" command at the destination volume.

### Syslog Message

Creating a common Snapshot copy for source volume "%s" and destination volume "%s" with relationship UUID "%s" has failed due to the following reason: "%s". Elapsed time since the latest successful common Snapshot copy is %s.

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

**error\_msg** (STRING): Error that prevented the common Snapshot copy from being created.

**css\_fail\_interval** (STRING): Interval between the last successful common Snapshot copy creation and recent failure to create a common Snapshot copy.

# sms.failed events

## sms.failed.precheck

### Severity

ERROR

### Description

This message occurs when a SnapMirror® Sync relationship cannot be established due to precheck failures.

### Corrective Action

Monitor the status of the SnapMirror Sync relationship by using the "snapmirror show" command. Retry the SnapMirror Sync operation after resolving the issue that caused the precheck failure.

### Syslog Message

SnapMirror Sync operation failed for source volume "%s" and destination volume "%s" with relationship UUID "%s" due to the following reason: "%s".

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.  
**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.  
**relationship\_id** (STRING): Relationship UUID.  
**error\_msg** (STRING): Reason for the failure.

# sms.fanout events

## sms.fanout.comm.snap.deleted

### Severity

ALERT

### Description

This message occurs when an older Snapshot(tm) copy is deleted as part of a SnapMirror® Synchronous resynchronize or update (common Snapshot copy) operation, which could lead to a "no common Snapshot scenario" between the synchronous and asynchronous disaster recovery (DR) copies that share the same source volume. If there is no common Snapshot copy between the synchronous and asynchronous DR copies, then a re-baseline will need to be performed during a disaster recovery.

### Corrective Action

You can ignore this message if there is no asynchronous relationship configured for the synchronous source volume. If there is an asynchronous relationship configured, then update the asynchronous relationship by using the "snapmirror update" command. The SnapMirror update operation will transfer the Snapshot copies that will act as common Snapshot copies between the synchronous and asynchronous destinations.

### Syslog Message

SnapMirror Synchronous operation '%s' for relationship '%s' has cleaned up some of the old base Snapshot copies between the synchronous source and synchronous destination, which could result in no common Snapshot copy existing between the synchronous and asynchronous destinations.

## Parameters

**sm\_operation** (STRING): SnapMirror Synchronous operation.  
**relationship\_id** (STRING): Relationship UUID.

# sms.giveback events

## sms.giveback.veto.override

### Severity

NOTICE

### Description

This message occurs during an SFO aggregate giveback, when the SnapMirror® subsystem indicates that giveback should be vetoed but the veto was overridden by the automated nondisruptive update procedure. The automated nondisruptive update procedure verifies the expected state of the aggregate.

### Corrective Action

(None).

## Syslog Message

snapmirror subsystem veto was overridden during giveback operation of "%s" aggregate.

## Parameters

**aggregate** (STRING): Name of the aggregate.

# sms.operation events

## sms.operation.blacklisted

### Severity

ERROR

### Description

This message occurs when a blacklisted operation is performed on the primary volume of a SnapMirror® Sync relationship when the relationship status is "Transitioning".

### Corrective Action

Auto-resync will attempt to bring the SnapMirror Sync relationship status to "InSync". You can also use the "snapmirror resync" command to change the relationship status. Avoid issuing blacklisted operations while the relationship status is "Transitioning".

## Syslog Message

SnapMirror Sync operation encountered a failure while the relationship status is "Transitioning" for relationship UUID "%s" because blacklisted operation "%s" was performed in the "%s" state.

## Parameters

**relationship\_id** (STRING): Relationship UUID.

**operation** (STRING): Blacklisted operation that was executed during the transitioning state.

**state** (STRING): Transitioning phase during which blacklisted operations are performed.

# sms.qr events

## sms.qr.mem.alloc.issue

### Severity

ERROR

### Description

This message occurs when waf1 fails to allocate required memory for storing streams attributes information or asked buffer memory length is too less to hold all the streams attributes info, in both the cases the status of Snapmirror Synchronous relationship changes to "OutOfSync".

### Corrective Action

Monitor the SnapMirror® Synchronous relationship status using the "snapmirror show" command. If the relationship status is not "InSync", resynchronize the relationship using "snapmirror resync" command.

## Syslog Message

Waf1 failed to allocate sufficient memory to hold all the streams attributes information for Snapmirror Synchronous relationship UUID: "%s" and File ID: "%s". The status of the Snapmirror Synchronous

relationship will be changed to "OutOfSync".

### Parameters

**relationship\_id** (STRING): Snapmirror Relationship UUID  
**file\_id** (STRING): File id for which the waf operation failed.

## sms.qr.ops.limit.exceed

### Severity

ERROR

### Description

This message occurs when SnapMirror® Synchronous detects that the number of write failures on the source volume of a SnapMirror® Synchronous relationship has exceeded the threshold. This can lead to change of the relationship status from "InSync" to "OutOfSync".

### Corrective Action

To recover from this situation, turn on sequential splitting in the Snapmirror® Synchronous relationship. For further information about correcting the problem, search the knowledgebase of the NetApp technical support web site for the "sms.qr.ops.limit.exceed" keyword.

### Syslog Message

The number of write failures on the source volume with Dsid: "%s" of the SnapMirror® Synchronous relationship UUID: "%s" has exceeded the threshold. This may lead to change of the relationship status from "InSync" to "OutOfSync".

### Parameters

**source\_volume\_dsid** (STRING): Source Volume DSID  
**relationship\_id** (STRING): Snapmirror Relationship UUID

## sms.qr.streams.limit.exceed

### Severity

ERROR

### Description

This message occurs when the number of streams associated with a file exceeds the maximum supported number of streams by SnapMirror® Synchronous. This will change the status of the Snapmirror Synchronous relationship to "OutOfSync".

### Corrective Action

Turn on sequential splitting in the SnapMirror Synchronous relationship. For further information about correcting the problem, search the knowledgebase of the NetApp technical support web site for the "sms.qr.streams.limit.exceed" keyword.

### Syslog Message

Maximum number of streams for a file supported by SnapMirror Synchronous relationship exceeded for relationship UUID: "%s" and file ID: "%s". Status of the relationship will be changed to "OutOfSync".

### Parameters

**relationship\_id** (STRING): SnapMirror relationship UUID.  
**file\_id** (STRING): File ID of the file for which the number of streams exceeded the maximum.



# sms.resync events

## sms.resync.attempt.failed

### Severity

ERROR

### Description

This message occurs when a resynchronize operation between the source volume and destination volume fails. The SnapMirror® Sync relationship is in "out-of-sync" status. Data protection is impacted.

### Corrective Action

Monitor SnapMirror Sync status using the "snapmirror show" command. If the auto-resync attempts fail, bring the relationship back to "in-sync" status manually by using the "snapmirror resync" command.

### Syslog Message

Resynchronize operation between source volume "%s" and destination volume "%s" with relationship UUID "%s" has failed. The next auto-resync will be attempted after "%d" mins.

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

**next\_resync\_interval** (INT): Interval after which the next resync will be attempted.

# sms.snap events

## sms.snap.not.replicated

### Severity

NOTICE

### Description

This message occurs when a Snapshot(tm) copy for SnapMirror® Synchronous relationship is not successfully replicated.

### Corrective Action

(None).

### Syslog Message

Snapshot copy '%s' is not successfully replicated for the relationship '%s' with source volume DSID '%llu' and path '%s'. Reason: %s.

### Parameters

**snapshot** (STRING): Name of the Snapshot copy to be replicated.

**transferId** (STRING): UUID of the SnapMirror transfer.

**volumeDSID** (LONGINT): Volume data set ID (DSID).

**volumePath** (STRING): Path of the volume.

**failureReason** (STRING): Reason for the failure.

## sms.src events

### sms.src.clear.fence.complete

#### Severity

NOTICE

#### Description

This message occurs when a SnapMirror® Synchronous operation has successfully cleared the persistent fence (read/write I/O has resumed) on the source volume.

#### Corrective Action

(None).

#### Syslog Message

Cleared persistent fence (read/write I/O has resumed) on the source volume for relationship UUID "%s" with source volume "%s" and destination volume "%s".

#### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

### sms.src.set.fence.complete

#### Severity

NOTICE

#### Description

This message occurs when a SnapMirror® Synchronous operation has successfully set the persistent fence (prevent read/write I/O) on the source volume.

#### Corrective Action

(None).

#### Syslog Message

Set persistent fence (prevent read/write I/O) on the source volume for relationship UUID "%s" with source volume "%s" and destination volume "%s".

#### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

## sms.status events

### sms.status.in.sync

#### Severity

NOTICE

## Description

This message occurs when a SnapMirror® Sync initialize or resynchronize operation is completed successfully, and the source volume and destination volume are in "in-sync" status.

## Corrective Action

(None).

## Syslog Message

Initialize or resynchronize operation was successful. Source volume "%s" and destination volume "%s" with relationship UUID "%s" are in "in-sync" status.

## Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

## sms.status.out.of.sync

### Severity

ERROR

### Description

This message occurs when a SnapMirror® Sync relationship changes status from "in-sync" to "out-of-sync". I/O restrictions are imposed on the source volume based on the mode of replication. For relationships with policy type "strict-sync-mirror", client read or write access to volume is prevented. Data protection is impacted.

### Corrective Action

Check network connection between the source and destination volumes. Monitor the SnapMirror Sync relationship status by using the "snapmirror show" command. Auto-resync will attempt to bring the relationship back to "in-sync" status.

### Syslog Message

Source volume "%s" and destination volume "%s" with relationship UUID "%s" is in "out-of-sync" status due to the following reason: "%s".

### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

**error\_msg** (STRING): Error that caused the relationship to change to "out-of-sync" status.

## sms.status.out.of.sync.cg

### Severity

ERROR

### Description

This message occurs when a SnapMirror for Business Continuity (SMBC) relationship changes status from "in-sync" to "out-of-sync". Due to this RPO=0 data protection will be disrupted.

### Corrective Action

Check the network connection between the source and destination volumes. Monitor the SMBC relationship status by using the "snapmirror show" command on the destination, and by using the "snapmirror list-destinations" command on the source. Auto-resync will attempt to bring the relationship back to "in-sync" status. If the resync fails, verify that all the nodes in the cluster are in quorum and are healthy.

### Syslog Message

Source CG "%s" and destination CG "%s" with relationship UUID "%s" is in "out-of-sync" status. Reason: "%s".

### Parameters

**srccgpath** (STRING): Source Consistency Group (CG) path.

**dstcgpath** (STRING): Destination CG path.

**cg\_relationship\_id** (STRING): CG relationship UUID.

**error\_msg** (STRING): Error that caused the relationship to change to "out-of-sync" status.

## sms.transition events

### sms.transition.retry

#### Severity

NOTICE

#### Description

This message occurs when a SnapMirror® initialize or resynchronize operation takes a long time to bring the relationship "in-sync" due to intermittent failure. This is not an error yet so retry has triggered. If it fails on retry, the EMS "sms.resync.attempt.failed" will be raised.

#### Corrective Action

(None).

#### Syslog Message

Initialize/Resynchronize operation between source volume "%s" and destination volume "%s" with relationship UUID "%s" is taking a long time to bring the relationship InSync due to intermittent failure causing a retry.

#### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

# smtape events

## smtape.bkp events

### smtape.bkp.abort.write.err

#### Severity

ERROR

#### Description

This message occurs when an SMTape backup operation is aborted because a write to tape device fails.

#### Corrective Action

Make sure that the tape device does not have media errors and is writable. Check for error details in the log file at `/mroot/etc/log/backup`.

#### Syslog Message

SMTape backup session %d from %s to %s was aborted because write to output failed.

#### Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**srcvolpath** (STRING): SMTape source volume.

**tapepath** (STRING): Current destination device to which backup was performed.

### smtape.bkp.fail

#### Severity

ERROR

#### Description

This message occurs when an SMTape backup operation to tape fails.

#### Corrective Action

Check for error details in the log file at `/mroot/etc/log/backup`. Retry the command after resolving the error reported.

#### Syslog Message

SMTape backup session %d from %s to %s failed with error %s.

#### Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**srcvolpath** (STRING): SMTape source volume.

**tapepath** (STRING): Current destination device to which backup was performed.

**errorstring** (STRING): Error string returned by SMTape.

### smtape.bkp.success

#### Severity

NOTICE

## Description

This message occurs when an SMTape backup operation to tape succeeds.

## Corrective Action

(None).

## Syslog Message

SMTape backup session %d from %s to tape was successful. %s, %d tape(s) written.

## Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**srcvolpath** (STRING): SMTape source volume.

**stats** (STRING): Statistics for the SMTape backup.

**numtapes** (INT): Number of tapes used.

## smtape.bkp.tape.busy

### Severity

ERROR

### Description

This message occurs when an SMTape backup to tape fails because the tape device is busy.

### Corrective Action

Make sure that the tape device is not being used by another application and retry the operation when it is free. Otherwise, issue the command using a tape device that is not currently in use.

### Syslog Message

Tape device is busy, aborting SMTape backup session %d from %s to %s.

### Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**srcvolpath** (STRING): SMTape source volume.

**tapepath** (STRING): SMTape destination tape device.

## smtape.bkp.tape.change

### Severity

NOTICE

### Description

This message occurs when a CLI-based SMTape backup operation to tape reaches end of tape.

### Corrective Action

Load the next tape media onto a drive and position it correctly, and then issue the 'system smtape continue -session-id [session-id] -tape [tapepath]' command with the new tape path.

### Syslog Message

SMTape backup session %d from %s to %s is waiting for a tape change.

## Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**srcvolpath** (STRING): SMTape source volume.

**tapepath** (STRING): Current destination tape device to which backup was performed.

# smtape.rst events

## smtape.rst.abort.chksum.err

### Severity

ERROR

### Description

This message occurs when an SMTape restore operation from tape is aborted due to checksum error.

### Corrective Action

Make sure that the correct tape is loaded and the tape is positioned correctly. If the tape is corrupted, contact NetApp technical support for the best effort recovery mechanisms.

### Syslog Message

SMTape restore session %d to %s from %s was aborted due to checksum error.

## Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**dstvolpath** (STRING): SMTape destination volume.

**tapepath** (STRING): Current source device from which restore was performed.

## smtape.rst.abort.read.err

### Severity

ERROR

### Description

This messages occurs when an SMTape restore operation is aborted because a read from the tape device fails.

### Corrective Action

Make sure that the input tape device and media are readable. Check for error details in the log file at /mroot/etc/log/backup.

### Syslog Message

SMTape restore session %d to %s from %s was aborted because read from input failed.

## Parameters

**sessionid** (INT): Session ID of the SMTape operation.

**dstvolpath** (STRING): SMTape destination volume.

**tapepath** (STRING): Current source device from which restore was performed.

## **smtape.rst.fail**

### **Severity**

ERROR

### **Description**

This message occurs when an SMTape restore operation fails.

### **Corrective Action**

Check for error details in the log file at /mroot/etc/log/backup. Retry the command after resolving the error reported.

### **Syslog Message**

SMTape restore session %d to %s from %s failed with error %s.

### **Parameters**

**sessionid** (INT): Session ID of the SMTape operation.

**dstvolpath** (STRING): SMTape destination volume.

**tapepath** (STRING): Current source device from which restore was performed.

**errorstring** (STRING): Error string returned by SMTape.

## **smtape.rst.success**

### **Severity**

NOTICE

### **Description**

This message occurs when an SMTape restore operation from tape succeeds.

### **Corrective Action**

(None).

### **Syslog Message**

SMTape restore session %d from tape to %s was successful. %s.

### **Parameters**

**sessionid** (INT): Session ID of the SMTape operation.

**dstvolpath** (STRING): SMTape destination volume.

**stats** (STRING): Statistics for the SMTape restore.

## **smtape.rst.tape.change**

### **Severity**

NOTICE

### **Description**

This message occurs when a CLI-based SMTape restore operation from tape reaches end of tape.

### **Corrective Action**

Load the next tape media onto a drive and position it correctly, and then issue the 'system smtape continue -session-id [session-id] -tape [tapepath]' command with the new tape path.



## **Syslog Message**

SMTape restore session %d from %s to %s is waiting for a tape change.

## **Parameters**

**sessionid** (INT): Session ID of the SMTape operation.

**tapepath** (STRING): Current source tape device from which restore was performed.

**dstvolpath** (STRING): SMTape destination volume.

# snapdiff events

## snapdiff.rebl events

### snapdiff.rebl.absent.allocated.block

#### Severity

INFORMATIONAL

#### Description

This message occurs when SnapDiff®, the snapshot copy differencing engine, encounters an absent allocated block. The absent allocated block is replicated to destination as raid error and Replication is unaffected by this issue.

#### Corrective Action

None. Zombie count of 0 is expected even when zombie processing is ongoing.

#### Syslog Message

Encountered an absent allocated block on inode %u in snapshot copy ID %u of volume %u Zombie processing: %u.

#### Parameters

**inode\_num** (INT): Identifier of the affected file.  
**snap\_id** (INT): Identifier of the Snapshot copy.  
**fsid** (INT): Identifier of the volume.  
**zombie\_processing** (INT): Ongoing zombie processing.

### snapdiff.rebl.BadAccessTime

#### Severity

NOTICE

#### Description

This message occurs when SnapDiff®, the snapshot copy differencing engine, encounters an invalid access time for an inode. Future accesses of the inode should not only succeed, but also reset its access time to a valid value. Replication is also unaffected by this issue.

#### Corrective Action

Run the (privilege: advanced) "inodepath -v volume\_name inode\_number" command to find the path of the affected file from the inode number, and modify the file within the Active File System. Note that file attributes captured in a snapshot copy cannot be changed.

#### Syslog Message

Encountered invalid access time on inode %u in Snapshot copy ID %u of volume %u.

#### Parameters

**inode\_num** (INT): Identifier of the affected file.  
**snap\_id** (INT): Identifier of the Snapshot copy.  
**fsid** (INT): Identifier of the volume.

## snapdiff.rebl.BadModifyTime

### Severity

NOTICE

### Description

This message occurs when SnapDiff®, the snapshot copy differencing engine, encounters an invalid modification time for an inode. Future accesses of the inode should not only succeed, but also reset its modification time to a valid value. Replication is also unaffected by this issue.

### Corrective Action

Run the (privilege: advanced) "inodepath -v volume\_name inode\_number" command to find the path of the affected file from the inode number, and modify the file within the Active File System. Note that file attributes captured in a snapshot copy cannot be changed.

### Syslog Message

Encountered invalid modify time on inode %u in Snapshot copy ID %u of volume %u.

### Parameters

**inode\_num** (INT): Identifier of the affected file.

**snap\_id** (INT): Identifier of the Snapshot copy.

**fsid** (INT): Identifier of the volume.

## snapdiff.spinnp events

### snapdiff.spinnp.abort.mgwd

### Severity

NOTICE

### Description

This message occurs when the system encounters an unexpected condition while computing differences between Snapshot(tm) copies using the SnapDiff v2 API.

### Corrective Action

(None).

### Syslog Message

Encountered unexpected Data ONTAP error while computing differences between Snapshot copies using the SnapDiff v2 API. SnapDiff session ID %s, FlexGroup %s, in vserver %s, with MSID %lu, diff Snapshot copy %s, diff Snapshot copy ID %lu, base Snapshot copy %s, base Snapshot copy ID %lu, current constituent DSID %lu. Error: %s

### Parameters

**sessionId** (STRING): SnapDiff v2 session ID.

**fgName** (STRING): FlexGroup name.

**vsName** (STRING): Vserver name.

**fgMsid** (LONGINT): FlexGroup MSID.

**diffSnapName** (STRING): Diff Snapshot copy name.

**diffSnapId** (LONGINT): Diff Snapshot copy logical ID.

**baseSnapName** (STRING): Base Snapshot copy name.

**baseSnapId** (LONGINT): Base Snapshot copy logical ID.

**curDsid** (LONGINT): Current constituent DSID.

**error** (STRING): Error message.

## snapdiff.zapi events

### snapdiff.zapi.filenameUnavailable

#### Severity

ERROR

#### Description

This message occurs when the Snapshot(tm) copy differencing engine, SnapDiff, fails to obtain an appropriate file name for the requested file access protocol. The SnapDiff session is aborted and future attempts will fail until the problem is resolved.

#### Corrective Action

(Call support).

#### Syslog Message

Failed to get an appropriate file name for the requested file access protocol for directory entry with inode %u in directory with inode %u in Snapshot copy %u of volume %u.

#### Parameters

**inode\_num** (INT): Identifier of the directory entry.

**parent\_inode\_num** (INT): Identifier of the directory.

**incr\_snap\_id** (INT): Identifier of the incremental Snapshot copy.

**fsid** (INT): Identifier of the volume.

**base\_snap\_id** (INT): Identifier of the base Snapshot copy.

# snaplock events

## snaplock.disk events

### snaplock.disk.import.7mode

#### Severity

ERROR

#### Description

This message occurs when the configuration management detects that an older(7-Mode) version of a SnapLock® disk is attached to a system running a Data ONTAP® release that cannot support it.

#### Corrective Action

Remove the SnapLock® disks from the system.

#### Syslog Message

File system %s has SnapLock® disks. This version of Data ONTAP® does not support older(7-Mode) SnapLock® versions

#### Parameters

**disk\_info** (STRING): Formatted information about the disk.

**shelf** (STRING): Shelf identifier where the disk is located.

**bay** (STRING): Disk bay within the shelf where disk is located.

**vendor** (STRING): Name of the vendor of the disk.

**model** (STRING): Model string of the disk drive.

**firmware\_revision** (STRING): Firmware revision number of the disk.

**serialno** (STRING): Serial number of the disk.

**disk\_type** (INT): Type of disk drive.

**disk\_rpm** (STRING): Rotational speed of disk in RPM.

**failure\_reason** (STRING): A text description of storage error leading to disk failure.

**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.

**site** (STRING): For a MetroCluster® configuration, indicates the site ("Local" or "Remote") where the disk is located. For non-MetroCluster configurations, the site is "Local".

### snaplock.disk.on.unsupported.version

#### Severity

ERROR

#### Description

This event is issued when configuration management detects that a SnapLock® disk is attached to a system running ONTAP® release which doesn't support SnapLock®.

#### Corrective Action

Install a version of Data ONTAP that supports SnapLock, or remove the SnapLock disks.

#### Syslog Message

File system %s has SnapLock disks. This version of Data ONTAP does not support SnapLock.

## Parameters

**disk\_info** (STRING): Formatted information of the disk  
**shelf** (STRING): Shelf identifier where the disk is located  
**bay** (STRING): Disk bay within the shelf where disk is located  
**vendor** (STRING): Name of the vendor of the disk  
**model** (STRING): Model string of the disk drive  
**firmware\_revision** (STRING): Firmware revision number of the disk  
**serialno** (STRING): Serial number of the disk  
**disk\_type** (INT): Type of disk drive  
**disk\_rpm** (STRING): Rotational speed of disk in RPM  
**failure\_reason** (STRING): A text description of storage error leading to disk failure.  
**carrier** (STRING): Unique ID of the carrier in which the disk drive is installed.  
**site** (STRING): For a MetroCluster configuration(tm), indicates the site {Local|Remote} where the disk is located. For non-MetroCluster configurations, site is 'Local'.

# snaplock.event events

## snaplock.event.based.retention.op.status

### Severity

NOTICE

### Description

This message occurs on completion of a event based retention apply operation on a SnapLock® volume.

### Corrective Action

(None).

### Syslog Message

An %s operation with ID %lu ended for policy "%s" and retention period "%s" on path %s in volume %s:%s with status: "%s", details: "%s". The operation passed on %lu files, failed on %lu files, skipped %lu files and ignored %lu inodes.

## Parameters

**operation\_type** (STRING): Type of the event based retention operation.  
**operation\_id** (LONGINT): Operation ID of the event based retention operation.  
**policy\_name** (STRING): Name of the event based retention policy.  
**retention\_period** (STRING): Retention period of the event based retention policy.  
**path** (STRING): Path of the form '/(directory path)' relative to volume root.  
**vserver** (STRING): Name of the vserver.  
**volume** (STRING): Name of the SnapLock volume, on which event based retention operation is running.  
**operation\_status** (STRING): Status of the event based retention operation.  
**status\_details** (STRING): Completion string of the event based retention operation.  
**num\_files\_processed** (LONGINT): Total count of files for which event based retention operation was performed.  
**num\_files\_failed** (LONGINT): Total count of files for which event based retention operation failed.  
**num\_files\_skipped** (LONGINT): Total count of files for which event based retention operation was skipped.  
**num\_inodes\_ignored** (LONGINT): Total count of inodes which were ignored by event based retention operation because they are not regular files.

# snaplock.invalid events

## snaplock.invalid.autocommit

### Severity

ERROR

### Description

This message occurs when an invalid autocommit period value is found on a SnapLock® volume. ONTAP® 9.4 software and later releases have a maximum autocommit period value of 10 years. If this maximum is exceeded, autocommit is disabled on the volume.

### Corrective Action

Correct the autocommit period on the SnapLock volume by using the "volume snaplock modify" command or the "volume-set-snaplock-attrs" Zephyr API.

### Syslog Message

Disabling autocommit on SnapLock volume %s%s on the Vserver with ID %s because the autocommit period exceeded the maximum of 10 years.

### Parameters

**volume** (STRING): Name of the SnapLock volume with invalid autocommit period value.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

# snaplock.legal events

## snaplock.legal.hold.dump.files.op.status

### Severity

NOTICE

### Description

This message occurs on completion of a legal-hold dump-files operation on a SnapLock® compliance volume.

### Corrective Action

(None).

### Syslog Message

A legal-hold dump-files operation with ID %lu ended for litigation %s on volume %s:%s and output path: "%s", with status: "%s", details: "%s". The operation passed on %lu files, failed on %lu files, skipped %lu files and ignored %lu inodes.

### Parameters

**operation\_id** (LONGINT): Operation ID of the legal-hold dump-files operation.

**litigation\_name** (STRING): Name of the litigation.

**vserver** (STRING): Name of the vserver.

**volume** (STRING): Name of the SnapLock compliance volume, on which legal-hold dump-files operation is running.

**output\_path** (STRING): Output path of legal-hold dump-files operation.

**operation\_status** (STRING): Status of the legal-hold dump-files operation.  
**status\_details** (STRING): Completion string of the legal-hold dump-files operation.  
**num\_files\_processed** (LONGINT): Total count of files for which legal-hold dump-files operation was performed.  
**num\_files\_failed** (LONGINT): Total count of files for which legal-hold dump-files operation failed.  
**num\_files\_skipped** (LONGINT): Total count of files for which legal-hold dump-files operation was skipped.  
**num\_inodes\_ignored** (LONGINT): Total count of inodes which were ignored by legal-hold dump-files operation because they are not regular files.

## snaplock.legal.hold.dump.litigations.job.status

### Severity

NOTICE

### Description

This message occurs on completion of a legal-hold dump-litigations operation.

### Corrective Action

(None).

### Syslog Message

A legal-hold dump-litigations operation with ID %s ended for %s with status: "%s", details: "%s".

### Parameters

**job\_id** (STRING): Job ID of the legal-hold operation.  
**vserver\_and\_volume** (STRING): Name of the vserver and volume (if any).  
**job\_status** (STRING): Status of the legal-hold dump-litigations job.  
**status\_details** (STRING): Completion string of the legal-hold dump-litigations job.

## snaplock.legal.hold.op.status

### Severity

NOTICE

### Description

This message occurs on completion of a legal-hold begin or end operation on a SnapLock® compliance volume.

### Corrective Action

(None).

### Syslog Message

A legal-hold %s operation with ID %lu ended for litigation %s on path %s in volume %s:%s with status: "%s", details: "%s". The operation passed on %lu files, failed on %lu files, skipped %lu files and ignored %lu inodes.

### Parameters

**operation\_type** (STRING): Type of the legal-hold operation.  
**operation\_id** (LONGINT): Operation ID of the legal-hold operation.  
**litigation\_name** (STRING): Name of the litigation.  
**path** (STRING): Path of the form '(directory path)' relative to volume root.  
**vserver** (STRING): Name of the vserver.



**volume** (STRING): Name of the SnapLock compliance volume, on which legal-hold operation is running.  
**operation\_status** (STRING): Status of the legal-hold operation.  
**status\_details** (STRING): Completion string of the legal-hold operation.  
**num\_files\_processed** (LONGINT): Total count of files for which legal-hold operation was performed.  
**num\_files\_failed** (LONGINT): Total count of files for which legal-hold operation failed.  
**num\_files\_skipped** (LONGINT): Total count of files for which legal-hold operation was skipped.  
**num\_inodes\_ignored** (LONGINT): Total count of inodes which were ignored by legal-hold operation because they are not regular files.

## snaplock.mcc events

### snaplock.mcc.veto.worm

#### Severity

ERROR

#### Description

This message occurs when SnapLock legal-hold or event based retention operations are aborted during a MetroCluster Switchback or a MetroCluster Switchover operation.

#### Corrective Action

(None).

#### Syslog Message

All SnapLock legal-hold and event based retention operations were aborted because a %s is in progress.

#### Parameters

**mcc\_op** (STRING): Describes the operation i.e. a MetroCluster Switchback or a MetroCluster Switchover.

## snaplock.nse events

### snaplock.NSE.conflict

#### Severity

NOTICE

#### Description

This message occurs when any of the following commands could not be run against a SnapLock Compliance disk, since doing so could risk the destruction of compliance data: "storage encryption disk sanitize", "storage encryption disk destroy", "storage encryption disk revert-to-original-state".

#### Corrective Action

(None).

#### Syslog Message

"storage encryption disk sanitize/destroy/revert-to-original-state" command could not be run on SnapLock Compliance disk '%s' as it could risk the destruction of compliance data.

#### Parameters

**disk** (STRING): Name of the SnapLock Compliance disk.

# snaplock.sfo events

## snaplock.sfo.veto.worm

### Severity

ERROR

### Description

This message occurs when SnapLock legal-hold or event-retention operations are aborted during an aggregate relocation or a sendhome operation.

### Corrective Action

(None).

### Syslog Message

All SnapLock legal-hold and event based retention operations were aborted for aggregate %s because %s is in progress.

### Parameters

**aggr\_name** (STRING): Name of the aggregate for which operations were aborted.

**arl\_or\_sendhome** (STRING): Describes the operation i.e. an aggregate relocation or a sendhome.

# snaplock.ss events

## snaplock.ss.lockBadVolume

### Severity

ALERT

### Description

This event is generated if an attempt is made to lock a snapshot on a SnapLock volume that has previously failed to find or load its SnapLock snapshot state.

### Corrective Action

Run waflliron on the volume.

### Syslog Message

A snapshot on SnapLock volume '%s%s' on the Vserver with ID '%s', could not be locked because the SnapLock snapshot state could not be found or loaded for the volume.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the snapshot could not be locked.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.badMagic

### Severity

ALERT

## Description

This event is generated if the SnapLock snapshot state on a SnapLock volume is corrupt with an invalid magic file identifier. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

## Corrective Action

Run waffliron on this volume.

## Syslog Message

The SnapLock snapshot state is corrupt with an invalid magic file identifier on SnapLock volume '%s%s' on the Vserver with ID '%s'.

## Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state is corrupt with an invalid magic file identifier.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.badSize

### Severity

ALERT

### Description

This event is generated if the SnapLock snapshot state on a SnapLock volume is corrupt with an incorrect size. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run waffliron on this volume.

### Syslog Message

The SnapLock snapshot state is corrupt with an incorrect size on SnapLock volume '%s%s' on the Vserver with ID '%s'.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state is corrupt with an incorrect size.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.badState

### Severity

ERROR

### Description

This message occurs when SnapLock® Snapshot® copies on the volume are corrupted, which prevents updates on the lockvault from succeeding until this is corrected.

### Corrective Action

Run the "snapshot fix-snaplock-state" command on this volume.

### Syslog Message

The SnapLock snapshot state is corrupt on SnapLock volume '%s%s' on the Vserver with ID '%s'.

### Parameters

**volume** (STRING): Name of the SnapLock volume on which the SnapLock Snapshot state is corrupt.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.badVersion

### Severity

ALERT

### Description

This event is generated if the SnapLock snapshot state on a SnapLock volume has an invalid version number. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run wafiron on this volume.

### Syslog Message

The SnapLock snapshot state has an invalid version number on SnapLock volume '%s%s' on the Vserver with ID '%s'. The bad version number found was %d and the expected version was %d.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state has an invalid version number.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

**bad\_version** (INT): The invalid version number of the SnapLock snapshot state.

**expected\_version** (INT): The expected version number of the SnapLock snapshot state.

## snaplock.ss.state.createFail

### Severity

ALERT

### Description

This event is generated if SnapLock is unable to create the state for the SnapLock snapshot state on a SnapLock volume. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run wafiron on this volume.

### Syslog Message

Failed to create the SnapLock snapshot state on SnapLock volume '%s%s' on the Vserver with ID '%s'.

## Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state could not be created.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.dir.createFail

### Severity

ALERT

### Description

This event is generated if SnapLock is unable to create the SnapLock state directory for storage of SnapLock state on a SnapLock volume.

### Corrective Action

Run waffliron on this volume.

### Syslog Message

Failed to create the SnapLock state directory for storage of SnapLock state on SnapLock volume '%s%s' on the Vserver with ID '%s'.

## Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock state directory could not be created.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.loadFail

### Severity

ALERT

### Description

This event is generated if the SnapLock snapshot state on a SnapLock volume failed to load. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run waffliron on this volume.

### Syslog Message

The SnapLock snapshot state could not be loaded from SnapLock volume '%s%s' on the Vserver with ID '%s'.

## Parameters

**volume** (STRING): The name of the SnapLock volume from which the SnapLock snapshot state could not be loaded.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.notFound

### Severity

ALERT

### Description

This event is generated if the SnapLock snapshot state cannot be found on a SnapLock volume. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run wafiron on this volume.

### Syslog Message

The SnapLock snapshot state could not be found on SnapLock volume '%s%s' on the Vserver with ID '%s'.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state could not be found.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.upgradeFail

### Severity

ALERT

### Description

This event is generated if SnapLock is unable to upgrade the the SnapLock snapshot state on a SnapLock volume. This volume will not allow SnapLock locking or unlocking of any further snapshots until this erroneous volume state is corrected.

### Corrective Action

Run wafiron on this volume.

### Syslog Message

Failed to upgrade the SnapLock snapshot state to SnapLock volume '%s%s' on the Vserver with ID '%s'.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state could not be written.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.ss.state.writeFail

### Severity

ALERT

### Description

This event is generated if SnapLock is unable to write the the SnapLock snapshot state on a SnapLock volume. This volume will not allow SnapLock locking or unlocking of any further snapshots until this

erroneous volume state is corrected.

### Corrective Action

Run wafliiron on this volume.

### Syslog Message

Failed to write the SnapLock snapshot state to SnapLock volume '%s%s' on the Vserver with ID '%s'.

### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock snapshot state could not be written.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.state events

### snaplock.state.dir.destroyFail

#### Severity

NOTICE

#### Description

This event is generated if SnapLock is unable to destroy the SnapLock state for a SnapLock volume. This will not affect the operation of the reverted volume and does not require any corrective action.

#### Corrective Action

(None).

#### Syslog Message

Failed to destroy the SnapLock state for SnapLock volume '%s%s' on the Vserver with ID '%s'.

#### Parameters

**volume** (STRING): The name of the SnapLock volume on which the SnapLock state could not be destroyed.

**app** (STRING): Application Name.

**volident** (STRING): Unique identifier of the volume when the volume name itself is insufficient.

## snaplock.sys events

### snaplock.sys.compclock.set

#### Severity

NOTICE

#### Description

This message occurs when the system's compliance clock time and date is set.

#### Corrective Action

(None).

## Syslog Message

The compliance clock time of the system has been set to %s. Reason: %s.

## Parameters

**date** (STRING): Human-readable compliance clock date and time.

**reason** (STRING): Reason for initialization of the system's compliance clock.

## snaplock.sys.compclock.skew.fixed

### Severity

NOTICE

### Description

This message occurs when the SnapLock® system compliance clock time is skewed from the storage system's system clock by more than one day and skew is subsequently fixed.

### Corrective Action

(None).

## Syslog Message

The system compliance clock time (%s) was skewed from the system clock (%s) by more than one day. The system compliance clock time has been reset to the system clock time.

## Parameters

**system\_compliance\_clock\_time** (STRING): Human readable system compliance clock time.

**system\_time** (STRING): Human readable system time.

## snaplock.trans events

### snaplock.trans.autocommit

### Severity

NOTICE

### Description

This message occurs when autocommit period is adjusted on a Snaplock volume upon transition. When a 7 mode SnapLock volume is transitioned to Data ONTAP® 9.4 and later, if we find the autocommit period to be greater than 10 years, we change it to 10 years, which is the maximum supported value for autocommit period in Data ONTAP® 9.4.

### Corrective Action

(None).

## Syslog Message

Autocommit period on SnapLock volume %s%s on the Vserver with ID %s is changed from %s to 10 years upon transition.

## Parameters

**volume** (STRING): Name of the SnapLock volume on which autocommit period was adjusted upon transition.

**app** (STRING): Application Name.



**volident** (STRING): Unique identifier the volume in cases where the volume name itself is insufficient.  
**old\_autocommit\_period** (STRING): Autocommit period of the SnapLock volume in 7 mode.

## snaplock.user events

### snaplock.user

#### Severity

NOTICE

#### Description

This message occurs when the SnapLock® logging system fails to securely log either the addition or deletion of the compliance admin role to or from a user account. The actual operation might have succeeded.

#### Corrective Action

No corrective action is necessary. However, a secure log of the request might not exist and the SnapLock log volume should be checked for correct operation.

#### Syslog Message

There was a problem securely logging the %s of the "%s" role %s the user "%s" by "%s". Reason: %s

#### Parameters

**event** (STRING): Identifies whether the role was added or deleted.  
**role** (STRING): Name of the compliance admin role.  
**tofrom** (STRING): Describes whether the role is being added to or deleted from the user account.  
**username** (STRING): User name to or from which the role is being added or deleted.  
**whodidit** (STRING): Name of the user performing the operation.  
**reason** (STRING): Reason for the logging failure.

## snaplock.vcc events

### snaplock.vcc.assoc.mismatch

#### Severity

INFORMATIONAL

#### Description

This message occurs when the system detects mismatch in the SnapLock® volume compliance clock association with the system compliance clock. This can happen when the volume compliance clock's node ID is different from system's node ID or when volume compliance clock's scc ID is different from system compliance clock's ID or when both differ from each other.

#### Corrective Action

(None).

#### Syslog Message

Volume compliance clock association on volume %s does not match the system compliance clock at %s and its compliance clock is now reassociated with the system compliance clock.

## Parameters

**volume** (STRING): Name of the volume that had the compliance clock association mismatch.  
**date** (STRING): Human readable system compliance clock time at which the VCC-SCC association mismatch is detected.

# snaplock.vol events

## snaplock.vol.compclock.skewed

### Severity

NOTICE

### Description

This event is generated when the SnapLock® volume compliance clock time is skewed from the SnapLock system compliance clock by more than one day.

### Corrective Action

(None).

### Syslog Message

The compliance clock time (%s) of volume '%s' is skewed from the system compliance clock (%s) by more than one day.

## Parameters

**volume\_compliance\_clock\_time** (STRING): Human readable volume compliance clock time.  
**volume\_name** (STRING): Volume name.  
**system\_compliance\_clock\_time** (STRING): Human readable system compliance clock time.

# snaplock.volume events

## snaplock.volume.repair

### Severity

NOTICE

### Description

This message occurs when wafiron is run on a SnapLock® volume. If inconsistency is detected, wafiron will repair SnapLock volume. File system data before inconsistency occurred and the file system data after repair, may be different. WORM protected data may be changed. Administrator must ensure that the data on the volume is verified and restored, if needed, from backup copy or replica.

### Corrective Action

(None).

### Syslog Message

%s repair such as wafiron has been issued on SnapLock %s '%s'. %s data may be changed.

## Parameters

**vtype1** (STRING): type of volume.  
**vtype2** (STRING): type of volume.

**volume** (STRING): The name of the SnapLock volume on which volume repair is issued.  
**vtype3** (STRING): type of volume.

# snapmirror events

## snapmirror.block events

### snapmirror.block.on.reconstruct

#### Severity

NOTICE

#### Description

This message occurs when a SnapMirror® operation needs to reconstruct the block type bits. The SnapMirror transfer is paused, resuming only after the block type bits for the data to be transferred have been reconstructed.

#### Corrective Action

If you do not want the SnapMirror transfer to pause until the block type bit reconstruction is complete, abort the transfer, allow the reconstruction to finish, and then try the SnapMirror transfer again.

#### Syslog Message

Starting reconstruction on %s. SnapMirror transfer from %s to %s:%s is blocked until the reconstruction is complete. If you do not want the SnapMirror transfer to be paused until the block type bit reconstruction is complete, abort the transfer, allow the reconstruction to finish, and then try the SnapMirror transfer again.

#### Parameters

**volume** (STRING): SnapMirror volume.  
**srcpath** (STRING): Replication source path.  
**dstfiler** (STRING): Replication destination system.  
**dstpath** (STRING): Replication destination path.

### snapmirror.block.reconstructErr

#### Severity

ERROR

#### Description

This event is issued when a SnapMirror operation needs to reconstruct block type bits, but this operation cannot be started.

#### Corrective Action

Please check for any other error message returned.

#### Syslog Message

Unable to start block type bits reconstruction on %s, %s

#### Parameters

**volume** (STRING): The snapmirror volume  
**error** (STRING): Description of the error.

# snapmirror.conf events

## snapmirror.conf.depre.cpsync

### Severity

ERROR

### Description

This message occurs when there an entry in the snapmirror.conf file is configured to use the outstanding value of greater than or equal to 10 seconds. This method of specifying cp-sync mode for the SnapMirror®, is being deprecated in Data ONTAP® 7.3. You should use the new keyword "semi-sync" to configure cp-sync mode. See the na\_snapmirror.conf man page for details.

### Corrective Action

Edit the /etc/snapmirror.conf file to convert the relationship to use the "semi-sync" keyword. See the man page for details about usage.

### Syslog Message

Entry for SnapMirror relationship %s:%s to %s:%s in the /etc/snapmirror.conf is using a deprecated format to specify cp-sync mode. See the na\_snapmirror.conf man page for the new format.

### Parameters

**src\_filer** (STRING): Source system.  
**src\_path** (STRING): Source path.  
**dst\_filer** (STRING): Destination system.  
**dst\_path** (STRING): Destination path.

## snapmirror.conf.full

### Severity

ERROR

### Description

This message occurs when the /etc/snapmirror.conf file in the SnapMirror® destination has more than the maximum allowed SnapMirror relationships.

### Corrective Action

Reduce the number of SnapMirror relationships in the /etc/snapmirror.conf file to the maximum allowed limit.

### Syslog Message

Warning from the SnapMirror configuration parser, error in line number '%d' of snapmirror.conf file, '%s'

### Parameters

**lineno** (INT): Line number in the /etc/snapmirror.conf file.  
**msg** (STRING): Reason for the error.

## snapmirror.conf.invalidStr

**Severity**

ERROR

**Description**

This message occurs on a replication destination appliance when Data ONTAP encounters invalid string while collecting replication statistics.

**Corrective Action**

Check snapmirror.conf file for any typing mistakes. In case of a relationship line, path names are typically specified as "filer:volume" or "filer:/vol/volume/qtree". In case of a connection line check for an equal sign. Check snapmirror.conf man page for discussion of path specifications.

**Syslog Message**

Illegal string: %s.

**Parameters**

**snapconfstr** (STRING): The relationship or connection line.

**snapmirror.conf.obsolete.nvsync****Severity**

ERROR

**Description**

This message occurs when there exists an entry in the snapmirror.conf file that is configured to use the outstanding value of greater than 0 seconds and fewer than 10 seconds. This mode (nvlog semi-sync) is obsolete in Data ONTAP® 7.3. When it is configured in this mode, the relationship is converted automatically to a fully synchronous relationship.

**Corrective Action**

Edit the /etc/snapmirror.conf file to convert the relationship to either a fully synchronous relationship with "sync" parameter or a cp semi-synchronous relationship with "semi-sync" parameter.

**Syslog Message**

SnapMirror relationship %s:%s to %s:%s is being converted to a fully synchronous relationship since the functionality provided by setting the outstanding parameter to a value greater than 0 and less than 10s in the snapmirror.conf file is obsolete.

**Parameters**

**src\_filer** (STRING): Source system.

**src\_path** (STRING): Source path.

**dst\_filer** (STRING): Destination system.

**dst\_path** (STRING): Destination path.

**snapmirror.copy events****snapmirror.copy.baseLine.done****Severity**

INFORMATIONAL

**Description**

This message occurs when the baseline transfer of a SnapMirror® copy for a volume is completed and the volume comes online.

**Corrective Action**

(None).

**Syslog Message**

Baseline transfer for volume %s completed.

**Parameters**

**dstpath** (STRING): Destination volume.

**snapmirror.copy.completed****Severity**

INFORMATIONAL

**Description**

This message occurs when the final update of a SnapMirror® copy for a volume is successfully completed and the volume is ready for transition.

**Corrective Action**

(None).

**Syslog Message**

Volume %s is now ready for transition.

**Parameters**

**dstpath** (STRING): Destination volume.

**snapmirror.copy.FailedErr****Severity**

ERROR

**Description**

This message occurs when a SnapMirror® copy relation gets an error notification from the SnapMirror engine.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror copy for volume %s failed. Reason for failure: %s.

**Parameters**

**dstpath** (STRING): Destination volume.

**reason** (STRING): Reason of failure.

## snapmirror.copy.pause.srcNoNewData

### Severity

INFORMATIONAL

### Description

This message occurs when the source volume of a SnapMirror® copy relation is typed 'replica read only' and has no new data to transfer. Transfer to its destination is paused.

### Corrective Action

If new data has been populated on the source volume by a SnapMirror schedule or manual update, run the 'snapmirror copy resume' command to restart the paused transfer. If no new data is expected on source volume, run the 'snapmirror copy complete' command to complete the volume transition.

### Syslog Message

The source contains no new data; transfer to destination %s is paused.

### Parameters

**dstpath** (STRING): Destination volume.

## snapmirror.copy.readyFinal

### Severity

INFORMATIONAL

### Description

This message occurs when the update time taken by the latest transfer is smaller than the delta time. The volume is now ready for final update.

### Corrective Action

(None).

### Syslog Message

Destination volume %s is ready for final update.

### Parameters

**dstpath** (STRING): Destination volume.

## snapmirror.corrupted events

### snapmirror.corrupted.replay

### Severity

ALERT

### Description

This message occurs when metadata replay fails in the process of making a Synchronous SnapMirror® destination flexible volume writable. This failure occurs because the filesystem on disk is corrupt.



### Corrective Action

This error will not allow the destination volume to be made writable. Wafiron must be run on the destination to correct the filesystem consistency. Once the inconsistency is fixed, snapmirror break can be attempted again.

### Syslog Message

Synchronous SnapMirror cannot proceed with snapmirror break for volume '%s' in containing aggregate '%s' because the filesystem on disk is corrupt.

### Parameters

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume  
**dst\_aggr** (STRING): The name of the aggregate containing the destination volume

## snapmirror.dst events

### snapmirror.dst.abortOps

#### Severity

ERROR

#### Description

This event is issued when a SnapMirror incremental transfer to a destination volume is nearing completion and has to abort any dump, vol copy, snapmirror or snapvault operations from that volume.

#### Corrective Action

(None).

#### Syslog Message

Aborting any dump, vol copy, aggr copy, snapmirror or snapvault operations from %s; snapmirror incremental transfer to it is nearing completion.

#### Parameters

**dst** (STRING): SnapMirror destination volume

### snapmirror.dst.autoRestrictMsg

#### Severity

NOTICE

#### Description

This event is issued when a SnapMirror initialize is issued on a volume or aggregate which is offline.

#### Corrective Action

SnapMirror initialize is issued on a restricted volume, use "vol/aggr restrict" to restrict the destination volume/aggregate.

#### Syslog Message

SnapMirror automatically restricting destination %s %s (currently offline).

## Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination volume

## snapmirror.dst.cannotFind

### Severity

ERROR

### Description

This event is issued when a SnapMirror incremental transfer is issued on a volume which is offline, restricted or does not exist.

### Corrective Action

Please make sure you entered the destination volume name correctly.

### Syslog Message

Destination %s %s is offline, is restricted, or does not exist, SnapMirror operation terminated.

## Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination

## snapmirror.dst.chk.stats

### Severity

INFORMATIONAL

### Description

Statistical information reported by "snapmirror check".

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**dstpath** (STRING): The destination volume.

**statistics** (STRING): Statistical information in the format: block\_no - offset : src\_val, dst\_val block\_no denotes whether it is vbn1 or 2. offset is the byte from the start of the block. src\_val is the value of that offset at the src, while dst\_val is at the destination.

## snapmirror.dst.chkDegradedMirrorErr

### Severity

ERROR

### Description

This message occurs when a replication 'check' operation is issued on a SnapMirror® destination volume that is smgr-mirrored and its smgr-mirror state is mirror-degraded or resyncing.

## Corrective Action

Restore the smgr-mirror state to normal or break the existing smgr-mirror relationship and try again. If available, see the diagnostic man pages and the 'aggr' command man pages for details.

## Syslog Message

Destination volume '%s' is smgr-mirrored and mirror-degraded or resyncing. Cannot execute a replication 'check' operation.

## Parameters

**dstpath** (STRING): The destination volume.

## snapmirror.dst.chkDone

### Severity

NOTICE

### Description

This event is issued when the destination in a SnapMirror®/SnapVault(tm) relationship has finished a replication 'check' operation and no mismatches were found on the destination volume or qtree.

### Corrective Action

(None).

### Syslog Message

Replication 'check' from %s:%s to %s completed and no mismatches were found.

### Parameters

**srcfiler** (STRING): The source appliance.

**srcpath** (STRING): The source path.

**dstpath** (STRING): The destination path.

## snapmirror.dst.chkDoneErr

### Severity

ERROR

### Description

This event is issued when the destination in a SnapMirror®/SnapVault(tm) relationship has finished a replication 'check' operation and mismatches were found on the destination volume or qtree.

### Corrective Action

Run replication fix ('check -F') to fix the mismatches in the relationship.

### Syslog Message

Replication 'check' from %s:%s to %s found mismatches. Please run replication fix ('check -F') to fix these mismatches.

### Parameters

**srcfiler** (STRING): The source appliance.

**srcpath** (STRING): The source path.

**dstpath** (STRING): The destination path.

## snapmirror.dst.chkSnapInfo

### Severity

INFORMATIONAL

### Description

This message occurs at the destination appliance of a SnapMirror® 'check' operation to notify the appliance about which base Snapshot(tm) copy will be used for the operation.

### Corrective Action

(None).

### Syslog Message

Snapmirror 'check' from '%s':'%s' to '%s' will use '%s' as the base Snapshot copy. All Snapshot copies older than and including the base Snapshot copy on the source volume will be used.

### Parameters

**srcfiler** (STRING): Source appliance.  
**srcpath** (STRING): Source path.  
**dstpath** (STRING): Destination path.  
**snapname** (STRING): Name of the base snapshot.

## snapmirror.dst.chkSyncErr

### Severity

ERROR

### Description

This message occurs on the destination system when a SnapMirror® 'check' operation is run on a Synchronous SnapMirror destination volume that has not been quiesced.

### Corrective Action

Quiesce the destination volume by running the 'snapmirror quiesce' command before starting the SnapMirror 'check' operation.

### Syslog Message

Could not run a SnapMirror 'check' operation on destination volume '%s', a Synchronous SnapMirror destination volume that has not been quiesced.

### Parameters

**dstpath** (STRING): Destination volume.

## snapmirror.dst.chkVersionErr

### Severity

ERROR

### Description

This message occurs on the destination system of a SnapMirror® 'check' operation when the source system's Data ONTAP® version does not support SnapMirror 'check' on FlexVol® volumes.

## Corrective Action

(None).

## Syslog Message

Data ONTAP version at source '%s' does not support SnapMirror 'check' on flexible volumes. Cannot execute a SnapMirror 'check' operation.

## Parameters

**srcfiler** (STRING): Source system.

## snapmirror.dst.directoryOperationFailASCII

### Severity

ERROR

### Description

This event is issued when a SnapVault or Qtree SnapMirror destination encounters an error in a directory operation - typically creating an entry, or trying to move an entry to purgatory (for rename or delete), and the name entry consists only of ASCII characters.

### Corrective Action

It may be necessary to break and reinitialize the relationship. Contact NetApp technical support.

### Syslog Message

Logical replication destination directory operation failed: source inode %d, name="%s%s%s".

### Parameters

**srclnum** (INT): The inode number of the file (or directory, link, etc.) on the source system.

**n1** (STRING): The first of three name components. When combined, these will either be the path to the file on the destination system (usual case), or the NFS name and NFS alternate name of the directory entry (in those cases where we run into trouble building the full path).

**n2** (STRING): The second of three name components.

**n3** (STRING): The last of three name components.

## snapmirror.dst.directoryOperationFailIntl

### Severity

ERROR

### Description

This event is issued when a SnapVault or Qtree SnapMirror destination encounters an error in a directory operation - trying to create, rename, or delete an entry. This event can occur as a result of the conditions that cause burts 106216 and 110298.

### Corrective Action

If the destination volume's "vol lang" was changed after the baseline transfer, then the relationship may need to be broken and reinitialized. Please be aware that "vol lang" does not take effect on the volume until the filer is rebooted; thus, if a baseline transfer occurs after the setting was changed but before the filer was rebooted, then the "vol lang" actually changes after the baseline transfer takes place. If no change in the destination volume's "vol lang" has occurred, you may need to change the source volume's language to match that on the destination volume. For guidelines on this procedure, please contact NetApp technical support.

## Syslog Message

Logical replication destination directory operation failed: source inode %d, name="%s%s%s". If the destination volume's "vol lang" was changed after the baseline transfer, then the relationship may need to be broken and reinitialized.

## Parameters

**srcInum** (INT): The inode number of the file (or directory, link, etc.) on the source system.

**n1** (STRING): The first of three name components. When combined, these will either be the path to the file on the destination system (usual case), or the NFS name and NFS alternate name of the directory entry (in those cases where we run into trouble building the full path).

**n2** (STRING): The second of three name components.

**n3** (STRING): The last of three name components.

## snapmirror.dst.fixDone

### Severity

NOTICE

### Description

This event is issued when the destination in a SnapMirror®/SnapVault(tm) relationship has successfully finished a replication fix ('check -F') operation.

### Corrective Action

(None).

## Syslog Message

Replication fix ('check -F') from %s:%s to %s completed successfully.

## Parameters

**srcfiler** (STRING): The source appliance.

**srcpath** (STRING): The source path.

**dstpath** (STRING): The destination path.

## snapmirror.dst.fixUsingPriorChk

### Severity

NOTICE

### Description

This event is issued when a relationship fix ('check -F') operation issued on a SnapMirror® destination volume is leveraging a previously executed replication 'check' operation and thus does not require to recompute the mismatches.

### Corrective Action

(None).

## Syslog Message

Replication fix ('check -F') from %s:%s to %s is leveraging a previously executed replication 'check' operation and does not require re-computation of the mismatches for the relationship.

## Parameters

**srcfiler** (STRING): The source appliance.

**srcpath** (STRING): The source path.

**dstpath** (STRING): The destination path.

## snapmirror.dst.multipath.connErr

### Severity

ERROR

### Description

This event is issued when the SnapMirror destination filer is unable to setup a multipath connection with the source filer. This can happen in case of a network connection failure or due to the lack of resources on the destination filer.

### Corrective Action

Check network connectivity and also make sure that a maximum number of SnapMirror transfers are not going on.

### Syslog Message

SnapMirror is unable to setup a multipath connection to %s, resorting to the standard single TCP connection.

## Parameters

**dstpath** (STRING): SnapMirror destination

## snapmirror.dst.multipath.maxConnErr

### Severity

NOTICE

### Description

This event is issued when the SnapMirror destination filer has reached the maximum number of active SnapMirror multipath connections that it can support.

### Corrective Action

Wait for some active SnapMirror transfers to complete and free up the necessary resources; then retry this transfer.

### Syslog Message

SnapMirror has reached the maximum number of active multipath connections, unable to allocate resources for a multipath transfer to %s.

## Parameters

**dstpath** (STRING): SnapMirror destination

## snapmirror.dst.noBuffers

### Severity

ERROR

**Description**

This event is issued when no buffers are available on the SnapMirror destination.

**Corrective Action**

(None).

**Syslog Message**

No buffers available on SnapMirror destination.

**Parameters**

(None).

**snapmirror.dst.noSnapshots****Severity**

ERROR

**Description**

SnapMirror requires the presence of a common base snapshot on the source and the destination for doing an incremental transfer. This event is issued when a SnapMirror incremental transfer is issued on a volume or aggregate which has no snapshots, as such the transfer is aborted.

**Corrective Action**

Since there are no snapshots on the destination volume/aggregate, the only way to establish a SnapMirror is by doing a SnapMirror initialize.

**Syslog Message**

SnapMirror found no snapshots in destination %s %s, aborting transfer.

**Parameters**

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination

**snapmirror.dst.OnlineCompres****Severity**

ERROR

**Description**

This message occurs when SnapMirror® cannot bring the destination online after a transfer because compression is not supported in this version of Data ONTAP®.

**Corrective Action**

Install a version of Data ONTAP that supports compression, available from your vendor.

**Syslog Message**

SnapMirror could not bring destination volume %s online. Compression is not supported in this version of Data ONTAP.

**Parameters**

**dst** (STRING): SnapMirror destination.



## snapmirror.dst.OnlineErr

### Severity

ERROR

### Description

This event is issued when SnapMirror is not able to bring the destination online after a transfer.

### Corrective Action

(None).

### Syslog Message

SnapMirror not able to bring destination %s %s online, %s.

### Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): SnapMirror destination

**error** (STRING): The error message

## snapmirror.dst.onlineFailed

### Severity

ERROR

### Description

This event is issued when after SnapMirror transfer it is not able to bring the destination online.

### Corrective Action

(None).

### Syslog Message

Failed to bring destination %s %s online (%d,%d).

### Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination

**errcode1** (INT): Error code 1

**errcode2** (INT): Error code 2

## snapmirror.dst.OnlineNotComp

### Severity

ERROR

### Description

This message occurs when SnapMirror® cannot bring the destination online after a transfer because SnapLock® is not supported on this version of Data ONTAP®.

### Corrective Action

Install a version of Data ONTAP that is supported with SnapLock, available from your vendor.

## Syslog Message

SnapMirror could not bring destination %s %s online. SnapLock is not supported on this version of Data ONTAP.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**dst** (STRING): SnapMirror destination.

## snapmirror.dst.requestDenied

### Severity

ERROR

### Description

This event is issued when the destination filer requesting a SnapMirror transfer from the source filer is denied the transfer request.

### Corrective Action

Please check the error message on source.

## Syslog Message

SnapMirror transfer request for %s from %s denied, %s.

## Parameters

**dstpath** (STRING): The SnapMirror destination, can be volume or qtree

**srcfiler** (STRING): The SnapMirror source filer

**error** (STRING): The error message

## snapmirror.dst.restart.createErr

### Severity

ERROR

### Description

This event is issued when SnapMirror is not able to create files required for a restart

### Corrective Action

(None).

## Syslog Message

SnapMirror not able to create %s which is needed for restart.

## Parameters

**filepath** (STRING): The file SnapMirror is not able to create

## snapmirror.dst.restart.failed

### Severity

NOTICE

**Description**

This event is issued when SnapMirror restart negotiation fails.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror restart negotiation failed.

**Parameters**

(None).

**snapmirror.dst.restart.invalidChkpt****Severity**

NOTICE

**Description**

This event is issued when an invalid restart checkpoint is detected on the destination.

**Corrective Action**

(None).

**Syslog Message**

Restart checkpoint for destination volume %s is no longer valid, removing checkpoint.

**Parameters**

**dst** (STRING): The SnapMirror destination volume

**snapmirror.dst.restart.invalidSnaps****Severity**

ERROR

**Description**

This very rare message occurs on a replication destination appliance, when SnapMirror® detects that a transfer restarted from a checkpoint is attempting to transfer additional Snapshot copies that were not recorded in the checkpoint. In this case, SnapMirror aborts the transfer and clears the checkpoint to prevent corruption on the destination appliance.

**Corrective Action**

No action is required for a scheduled SnapMirror transfer. The next scheduled transfer starts from the beginning. If you manually initiated this transfer, manually start a new transfer.

**Syslog Message**

SnapMirror cannot restart the transfer from %s:%s to %s due to a change in the state of Snapshot copies on the source appliance. The checkpoint has been cleared, and the next transfer will start from beginning.

**Parameters**

**srcfiler** (STRING): Source appliance.

**srcpath** (STRING): SnapMirror source volume.

**dstpath** (STRING): SnapMirror destination volume.

## **snapmirror.dst.restart.noRootVol**

### **Severity**

ERROR

### **Description**

This event is issued when SnapMirror destination is not able to find the root volume.

### **Corrective Action**

(None).

### **Syslog Message**

Cannot find root volume.

### **Parameters**

(None).

## **snapmirror.dst.restart.readErr**

### **Severity**

ERROR

### **Description**

This event is issued when SnapMirror is not able to read files which were created for a restart.

### **Corrective Action**

(None).

### **Syslog Message**

SnapMirror not able to read %s which is needed for restart.

### **Parameters**

**filepath** (STRING): The file SnapMirror is not able to read

## **snapmirror.dst.restart.writeErr**

### **Severity**

ERROR

### **Description**

This event is issued when SnapMirror is not able to write to files which were created for a restart.

### **Corrective Action**

(None).

### **Syslog Message**

SnapMirror not able to write to %s which is needed for restart.

## Parameters

**filepath** (STRING): The file SnapMirror is not able to write to

## snapmirror.dst.resync.info

### Severity

NOTICE

### Description

This event is issued when a SnapMirror resync operation is successful and tells the base snapshot being used.

### Corrective Action

(None).

### Syslog Message

SnapMirror resync of %s to %s:%s is using %s as the base snapshot.

## Parameters

**dstpath** (STRING): SnapMirror destination volume or qtree

**srcfiler** (STRING): SnapMirror source filer

**srcpath** (STRING): SnapMirror source volume or qtree

**basesnapname** (STRING): The base snapshot being used.

## snapmirror.dst.resync.noAction

### Severity

NOTICE

### Description

This event is issued when a SnapMirror resync -n is executed.

### Corrective Action

(None).

### Syslog Message

SnapMirror resync of %s not performed because -n(ot really) option set.

## Parameters

**dstpath** (STRING): SnapMirror destination

## snapmirror.dst.revertAggr

### Severity

NOTICE

### Description

This event is issued when a SnapMirror initialize operation has to revert the checksum state of the destination aggregate because the source and destination aggregate have differing checksum states.

## Corrective Action

No corrective action required.

## Syslog Message

A SnapMirror transfer to destination aggregate %s has found that the source and destination aggregates have differing checksum states. Data ONTAP is aborting the transfer and reverting the checksum state for the destination aggregate. After Data ONTAP displays a message showing the revert is complete, you will need to retry the operation. You can use the "aggr status -c" command to see the progress of the reversion.

## Parameters

**dst** (STRING): The SnapMirror destination aggregate

## snapmirror.dst.revertErr

### Severity

ERROR

### Description

This event is issued when a SnapMirror initialize operation is not able to revert the checksum state of the destination volume.

### Corrective Action

(None).

## Syslog Message

Data ONTAP cannot revert the checksum state of destination %s %s.

## Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination

## snapmirror.dst.reverting

### Severity

ERROR

### Description

This event is issued when a SnapMirror initialize operation is issued on a volume/aggregate which is reverting its checksum state.

### Corrective Action

(None).

## Syslog Message

Aborting SnapMirror transfer to %s %s because it is reverting its checksum state.

## Parameters

**volaggr** (STRING): Whether volume or aggregate

**dst** (STRING): The SnapMirror destination volume

## snapmirror.dst.revertVol

### Severity

ERROR

### Description

This event is issued when a SnapMirror initialize operation has to revert the checksum state of the destination volume because the source and destination have differing checksum states.

### Corrective Action

No corrective action required.

### Syslog Message

A SnapMirror transfer to destination volume %s has found that the source and destination volumes have differing checksum states. Data ONTAP is aborting the transfer and reverting the checksum state for the destination volume. After Data ONTAP displays a message showing the revert is complete, you will need to retry the operation. You can use the "vol status -c" command to see the progress of the reversion.

### Parameters

**dst** (STRING): The SnapMirror destination

## snapmirror.dst.snapCreateWarning

### Severity

NOTICE

### Description

This event is issued when a named snapshot creation at the end of a qtree SnapMirror transfer fails. Since the transfer succeeded, only the snap create in the end failed, we just issue a warning.

### Corrective Action

(None).

### Syslog Message

SnapMirror transfer from %s:%s to %s : creation of named snapshot "%s" failed.

### Parameters

**srcfiler** (STRING): The SnapMirror source filer

**srcpath** (STRING): The SnapMirror source qtree

**dstpath** (STRING): The SnapMirror destination qtree

**snapname** (STRING): The snapshot which was supposed to be created on the destination volume

## snapmirror.dst.snapDelErr

### Severity

ERROR

### Description

This event is issued when on a SnapMirror incremental the snapshot that SnapMirror tries to delete in the destination volume is softlocked.

**Corrective Action**

(None).

**Syslog Message**

Snapshot %s in destination volume %s is in use, cannot delete.

**Parameters**

**snapname** (STRING): The snapshot which is locked

**dst** (STRING): The SnapMirror destination volume

**snapmirror.dst.sync.chkErr****Severity**

ERROR

**Description**

This message occurs when a SnapMirror® destination volume tries to go into sync while an active SnapMirror 'check' operation is in progress.

**Corrective Action**

The mirror gets back into sync after the SnapMirror 'check' operation is complete. To expedite this process, abort the ongoing SnapMirror 'check' operation by running the 'snapmirror check abort' command on the destination system.

**Syslog Message**

The destination volume '%s' has an active SnapMirror 'check' operation in progress. Could not go into sync.

**Parameters**

**dstpath** (STRING): Destination volume.

**snapmirror.dst.sync.mm.nvlog****Severity**

ERROR

**Description**

This message occurs when NVLOG version mismatch between the source and destination appliance prevents the use of synchronous SnapMirror®. The source and destination appliances must use the same version of NVLOG to mirror synchronously. Data ONTAP® will use periodic updates to mirror the volume until the mismatch is resolved.

**Corrective Action**

Upgrade the source or destination appliance so that they both use the same NVLOG version. It is also recommended that both appliances run the same version of Data ONTAP.

**Syslog Message**

NVLOG version mismatch for Synchronous SnapMirror from %s:%s to %s. The source and destination appliances must use the same NVLOG version to mirror synchronously. Reverting to periodic updates to mirror the volume.



## Parameters

**srcfiler** (STRING): Name of the SnapMirror source appliance.  
**srcpath** (STRING): Name of the SnapMirror source volume.  
**dstpath** (STRING): Name of the SnapMirror destination volume.

## snapmirror.dst.sync.mm.proto

### Severity

ERROR

### Description

This message occurs when a protocol version mismatch between the source and destination appliance prevents the use of synchronous SnapMirror®. The source and destination appliances must use the same version of protocol to mirror synchronously. Data ONTAP® will use periodic updates to mirror the volume until the mismatch is resolved.

### Corrective Action

Upgrade the source or destination appliance so that they both use the same protocol version. It is also recommended that both appliances run the same version of Data ONTAP.

### Syslog Message

Protocol version mismatch for Synchronous SnapMirror from %s:%s to %s. The source and destination appliances must use the same protocol version to mirror synchronously. Reverting to periodic updates to mirror the volume.

## Parameters

**srcfiler** (STRING): Name of the SnapMirror source appliance.  
**srcpath** (STRING): Name of the SnapMirror source volume.  
**dstpath** (STRING): Name of the SnapMirror destination volume.

## snapmirror.dst.updateDelayed

### Severity

NOTICE

### Description

This event is issued when an update request is received on a destination qtree, which is being used as the source for another transfer.

### Corrective Action

(None).

### Syslog Message

Transfer from %s to %s:%s is in progress, SnapMirror transfer from %s:%s delayed.

## Parameters

**dstpath** (STRING): The SnapMirror destination  
**dstfiler2** (STRING): The destination filer for dstpath  
**dstpath2** (STRING): The destination for dstpath  
**srcfiler** (STRING): The source filer for dstpath  
**srcpath** (STRING): The source for dstpath

# snapmirror.flextraderr events

## snapmirror.flexToTradErr

### Severity

ERROR

### Description

This message occurs when the user tries to copy a flexible volume into a traditional volume using SnapMirror®.

### Corrective Action

Retry the operation with a flexible volume as the destination.

### Syslog Message

snapmirror: A flexible volume cannot be copied into a traditional volume.

### Parameters

(None).

# snapmirror.log events

## snapmirror.log.openErr

### Severity

ERROR

### Description

This event is issued when an error is encountered when opening the SnapMirror log file for writing.

### Corrective Action

(None).

### Syslog Message

Failed to open SnapMirror log file: %s.

### Parameters

**error** (STRING): The error encountered while opening the SnapMirror log file for writing.

## snapmirror.log.unknownEvent

### Severity

ERROR

### Description

This event is issued when an unknown event is sent to be logged in the SnapMirror log file.

### Corrective Action

(None).

## Syslog Message

SnapMirror log encountered unknown event : %d.

## Parameters

**event** (INT): The unknown event

## snapmirror.log.writeErr

### Severity

ERROR

### Description

This event is issued when an error is encountered when writing to the SnapMirror log file.

### Corrective Action

(None).

## Syslog Message

Failed to write to SnapMirror log file: %s.

## Parameters

**error** (STRING): The error encountered while writing to the SnapMirror log file.

# snapmirror.lrs events

## snapmirror.lrs.FiletypeMod

### Severity

ALERT

### Description

This message occurs when SnapMirror encounters mismatched file types between two Snapshot copies that it is comparing. The transfer is aborted and future attempts fail until the problem is resolved.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

SnapMirror encountered mismatched file types for inode %u in volume %u. It has file type %u in Snapshot copy %u and a file type %u in Snapshot copy %u.

## Parameters

**inode\_num** (INT): Identifier of the affected file.

**fsid** (INT): FSID of the volume.

**base\_filetype** (INT): File type of the file in the base Snapshot copy.

**base\_snap\_id** (INT): Identifier of the base Snapshot copy.

**incr\_filetype** (INT): File type of the file in the incremental Snapshot copy.

**incr\_snap\_id** (INT): Identifier of the incremental Snapshot copy.

# snapmirror.metadata events

## snapmirror.metadata.create.fail

### Severity

ERROR

### Description

This message occurs when a Synchronous SnapMirror® destination flexible volume is unable to enter synchronous replication mode because the containing aggregate does not have enough space left to allow the update to proceed.

### Corrective Action

This error can be fixed by creating more space in the containing aggregate. This can be done by adding additional disks to the containing aggregate.

### Syslog Message

Synchronous SnapMirror failed to update volume '%s' because there is not enough space in containing aggregate '%s'.

### Parameters

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume  
**dst\_aggr** (STRING): The name of the aggregate containing the destination volume

## snapmirror.metadata.definite.loss

### Severity

ERROR

### Description

This message occurs when a Synchronous SnapMirror® destination flexible volume has lost metadata while entering synchronous replication mode.

### Corrective Action

This error becomes obsolete if the destination volume successfully goes into sync. If the destination volume is made writable before it goes into sync, then it will come online with a recovery point (a snapshot) that's behind the on-disk recovery point that the destination had when it fell out of sync. Wafiron and Snapmirror Checker must be run on the destination after it is made writable to verify filesystem consistency.

### Syslog Message

Synchronous SnapMirror has lost metadata for volume '%s' in containing aggregate '%s' when going into sync.

### Parameters

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume  
**dst\_aggr** (STRING): The name of the aggregate containing the destination volume

## snapmirror.metadata.possible.loss

**Severity**

ERROR

**Description**

This message occurs when a Synchronous SnapMirror® destination flexible volume may or may not have lost metadata while entering synchronous replication mode.

**Corrective Action**

This warning becomes obsolete if the destination volume successfully goes into sync. If the destination volume is made writable before it goes into sync, then there are two possible outcomes. If the destination comes online with a recovery point (a snapshot) that's behind the on-disk recovery point that the destination had when it fell out of sync, then the system has lost the recovery point. If however the destination comes online with a recovery point (a snapshot) that's better than before then there are no problems. In either case, wafilon and Snapmirror Checker must be run on the destination after it is made writable to check for any possible errors.

**Syslog Message**

Synchronous SnapMirror may or may not have lost metadata for volume '%s' in containing aggregate '%s' when going into sync.

**Parameters**

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume  
**dst\_aggr** (STRING): The name of the aggregate containing the destination volume

## snapmirror.migrate events

### snapmirror.migrate.complete

**Severity**

NOTICE

**Description**

This event is issued when a migrate completes successfully.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror migrate from %s:%s to %s:%s complete.

**Parameters**

**srcfiler** (STRING): SnapMirror source system.  
**srcpath** (STRING): SnapMirror source volume.  
**dstfiler** (STRING): SnapMirror destination system.  
**dstpath** (STRING): SnapMirror destination volume.

### snapmirror.migrate.dstErr

**Severity**

ERROR

## Description

This event is issued when a migrate request encounters an error on the SnapMirror® destination.

## Corrective Action

(None).

## Syslog Message

SnapMirror migration from %s:%s to %s:%s reports error: %s.

## Parameters

**srcfiler** (STRING): SnapMirror source system.

**srcpath** (STRING): SnapMirror source volume.

**dstfiler** (STRING): SnapMirror destination system.

**dstpath** (STRING): SnapMirror destination volume.

**error** (STRING): Error encountered on the SnapMirror destination.

## snapmirror.migrate.entryFound

### Severity

ERROR

### Description

This event is issued when partially completed SnapMirror migrate entry is encountered while parsing the registry at boot time looking for SnapMirror migrate entries.

### Corrective Action

(None).

### Syslog Message

SnapMirror found %s in mid-migration state : %s.

### Parameters

**volume** (STRING): The volume whose migrate entry is found

**error** (STRING): The error string

## snapmirror.migrate.fsidErr

### Severity

ERROR

### Description

This message occurs when a migrate request encounters an error after rewriting the File System Identifier (FSID) of a volume. During the migration process, the FSIDs of both source and destination sides are rewritten. When this process encounters an error, the migration process must roll back both volumes to their original states. This message notifies you that it cannot restore their FSIDs back to their original FSIDs.

### Corrective Action

Set the privilege level to advanced. Rewrite the FSID of the volume to its original FSID using the 'vol rewrite\_fsid' command.

## Syslog Message

Cannot rewrite volume %s FSID to its original FSID after migration failure: to 0x%x

## Parameters

**volume** (STRING): Volume name where the problem is detected.

**toFSID** (INT): Original FSID of the volume.

## snapmirror.migrate.incorrectEntry

### Severity

ERROR

### Description

This event is issued when an incorrect migrate entry is encountered while parsing the registry at boot time trying to find out partially completed SnapMirror migrate operations.

### Corrective Action

(None).

## Syslog Message

%s has an unrecognizable SnapMirror migrate state entry associated with it : %s.

## Parameters

**volume** (STRING): The volume whose migrate entry is found

**error** (STRING): The error string

## snapmirror.migrate.registryErr

### Severity

ERROR

### Description

This event is issued when an error is encountered parsing the registry at boot time when trying to find out partially completed SnapMirror migrate operations.

### Corrective Action

(None).

## Syslog Message

Unable to find SnapMirror migrate entries in the registry, registry parsing error : %s.

## Parameters

**error** (STRING): The error encountered while parsing the registry looking for partially completed migrate operations.

## snapmirror.migrate.srcErr

### Severity

ERROR

**Description**

This event is issued when a migrate request encounters an error on the SnapMirror® source.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror migration from %s:%s to %s:%s reports error: %s.

**Parameters**

**srcfiler** (STRING): SnapMirror source system.  
**srcpath** (STRING): SnapMirror source volume.  
**dstfiler** (STRING): SnapMirror destination system.  
**dstpath** (STRING): SnapMirror destination volume.  
**error** (STRING): Error encountered on the SnapMirror source.

**snapmirror.migrate.srcNotReady****Severity**

ERROR

**Description**

This event is issued when a migrate request is received for a source volume which is not yet ready.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror migrate request for volume %s, which is not in migrating state.

**Parameters**

**src** (STRING): SnapMirror source volume

**snapmirror.migrate.unknownState****Severity**

NOTICE

**Description**

This event is issued when an unknown state for a volume in migration is encountered.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror found volume in illegal migration state : %d.

**Parameters**

**state** (INT): Migration state of the volume



# snapmirror.multipath events

## snapmirror.multipath.initErr

### Severity

ERROR

### Description

This event is issued during boot or during takeover when the filer is unable to setup the necessary infrastructure required for SnapMirror multipathing support.

### Corrective Action

Check the network connectivity on the source and the destination side. If the problem persists then, contact NetApp technical support

### Syslog Message

SnapMirror is unable to setup multipath connection support on %s host.

### Parameters

**host** (STRING): The host, local or partner

# snapmirror.reconstruct events

## snapmirror.reconstruct.complete

### Severity

NOTICE

### Description

This message occurs when the SnapMirror® requested reconstruction of a per-volume meta file used for labeling block information is complete.

### Corrective Action

No corrective action.

### Syslog Message

Reconstruction %s for data to be transferred by SnapMirror from %s to %s:%s was completed successfully in %u secs.

### Parameters

**volume** (STRING): SnapMirror volume.  
**srcpath** (STRING): Replication source path.  
**dstfiler** (STRING): Replication destination system.  
**dstpath** (STRING): Replication destination path.  
**time** (INT): Time it took for the reconstruction to be complete.

# snapmirror.registry events

## snapmirror.registry.cleanup

### Severity

INFORMATIONAL

### Description

This message occurs when the system cleans up an invalid value in SnapMirror® registry.

### Corrective Action

(None).

### Syslog Message

Cleanup invalid registry %s: %s.

### Parameters

**type** (STRING): SnapMirror registry key or value .

**invalid** (STRING): Invalid SnapMirror registry key or value.

## snapmirror.replay events

### snapmirror.replay.failed

### Severity

EMERGENCY

### Description

This message occurs when metadata replay fails in the process of making a Synchronous SnapMirror® destination flexible volume writable. This will result in the destination volume coming online with the view of the last snapshot successfully transferred instead of the view of the active filesystem up until the last consistency point..

### Corrective Action

There is no way to recover from this data loss.

### Syslog Message

Synchronous SnapMirror failed to replay metadata up until the last consistency point during snapmirror break for volume '%s'.

### Parameters

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume

## snapmirror.resync events

### snapmirror.resync.blkRclmtn

### Severity

ERROR

### Description

This message occurs when a flexible volume restored from a Snapshot(tm) copy might not be able to

relinquish free blocks to its containing aggregate. As a result, the containing aggregate might run out of space if this condition continues past the next incremental SnapMirror® transfer.

### Corrective Action

Monitor the state after the next incremental SnapMirror update for this volume to verify that space is available in the containing aggregate. Free up space in the containing aggregate if the condition persists.

### Syslog Message

Flexible volume '%s' cannot release free blocks to its containing aggregate '%s'.

### Parameters

**vol** (STRING): Volume name.

**aggr** (STRING): Aggregate name.

## snapmirror.retry events

### snapmirror.retry.replay

#### Severity

ERROR

#### Description

This message occurs when metadata replay fails in the process of making a Synchronous SnapMirror® destination flexible volume writable. This failure occurs because sufficient space was not available in the containing aggregate.

#### Corrective Action

This error will not allow the destination volume to be made writable until additional space is made available in the containing aggregate. The snapmirror break command should be issued again after additional disks have been added to the aggregate.

#### Syslog Message

Synchronous SnapMirror cannot proceed with snapmirror break for volume '%s' in containing aggregate '%s' because sufficient space was not available.

#### Parameters

**dst\_volume** (STRING): The name of the Synchronous SnapMirror® destination volume

**dst\_aggr** (STRING): The name of the aggregate containing the destination volume

## snapmirror.sockreaderr events

### snapmirror.sockReadErr

#### Severity

ERROR

#### Description

This event is issued when SnapMirror tries to read a very large chunk from the network.

**Corrective Action**

(None).

**Syslog Message**

Attempted illegal network read of %d bytes.

**Parameters**

**bytes** (INT): Bytes SnapMirror tried to read from socket.

## snapmirror.sockstrreaderr events

**snapmirror.sockStrReadErr****Severity**

ERROR

**Description**

This event is issued when SnapMirror tries to read a large string from the socket.

**Corrective Action**

(None).

**Syslog Message**

Attempted read of a large string of length %d.

**Parameters**

**len** (INT): String length SnapMirror tried to read from socket.

## snapmirror.sockstrwriteerr events

**snapmirror.sockStrWriteErr****Severity**

ERROR

**Description**

This event is issued when SnapMirror tries to write a large string to the socket.

**Corrective Action**

(None).

**Syslog Message**

Write string too large %d.

**Parameters**

**len** (INT): String length SnapMirror tried to write to socket.

# snapmirror.src events

## snapmirror.src.cannotFind

### Severity

ERROR

### Description

This event is issued when a SnapMirror is issued from a source volume which is offline, restricted or does not exist.

### Corrective Action

(None).

### Syslog Message

SnapMirror : source %s %s is offline, is restricted, or does not exist.

### Parameters

**volaggr** (STRING): Whether volume or aggregate

**src** (STRING): The SnapMirror source

## snapmirror.src.chkDone

### Severity

NOTICE

### Description

This event is issued when the source in a SnapMirror@/SnapVault(tm) relationship has finished a replication 'check' operation. The user should review the destination appliance's SnapMirror log file to make sure that no mismatches were found.

### Corrective Action

Please review destination appliance's SnapMirror log file to make sure that no mismatches were found in the relationship.

### Syslog Message

Replication 'check' from %s to %s:%s completed: Please review destination appliance's SnapMirror log file for more information.

### Parameters

**srcpath** (STRING): The source path.

**dstfiler** (STRING): The destination appliance.

**dstpath** (STRING): The destination path.

## snapmirror.src.chkSnapInfo

### Severity

INFORMATIONAL

## Description

This message occurs at the source appliance of a SnapMirror® 'check' operation to notify the appliance about which base Snapshot(tm) copy will be used for the operation.

## Corrective Action

(None).

## Syslog Message

Snapmirror 'check' from '%s' to '%s!:%s' will use '%s' as the base Snapshot copy. All Snapshot copies older than and including the base Snapshot copy on the source volume will be used.

## Parameters

**srcpath** (STRING): Source path.  
**dstfiler** (STRING): Destination appliance.  
**dstpath** (STRING): Destination path.  
**snapname** (STRING): Name of the base snapshot.

## snapmirror.src.connDropped

### Severity

ERROR

### Description

This event is issued when SnapMirror source gets an error reading/writing to the network when communicating with the destination at the beginning of a SnapMirror transfer.

### Corrective Action

(None).

### Syslog Message

Error reading/writing to network, connection dropped.

### Parameters

(None).

## snapmirror.src.diffSnap

### Severity

ERROR

### Description

This event is issued when SnapMirror destination requests for a different snapshot than expected.

### Corrective Action

(None).

### Syslog Message

SnapMirror transfer from %s to %s:%s : snapshot %s confusion, aborting transfer.

### Parameters

**srcpath** (STRING): The SnapMirror source volume or qtree

**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination volume or qtree  
**snapname** (STRING): The snapshot requested by the destination

## snapmirror.src.fixDone

### Severity

NOTICE

### Description

This event is issued when the source in a SnapMirror®/SnapVault(tm) relationship has finished a replication fix ('check -F') operation. The user should check the destination appliance's SnapMirror log file to make sure that all fixes were applied successfully.

### Corrective Action

Please review destination appliance's SnapMirror log file to make sure that all mismatches were fixed for the relationship.

### Syslog Message

Replication fix ('check -F') from %s to %s:%s completed: Please review destination storage appliance's SnapMirror log file for more information.

### Parameters

**srcpath** (STRING): The source path.  
**dstfiler** (STRING): The destination appliance.  
**dstpath** (STRING): The destination path.

## snapmirror.src.inCutOver

### Severity

ERROR

### Description

This message occurs under two conditions: 1) When a new SnapMirror® command is issued on the destination volume of a volume being moved using the 'vol move' command. 2) When a new SnapMirror command is issued on the source volume of a volume being moved using the 'vol move' that is in the cutover phase. The new SnapMirror operation will abort.

### Corrective Action

Retry the new SnapMirror operation after the current 'vol move' operation is complete.

### Syslog Message

SnapMirror: source volume %s is part of a 'vol move' operation.

### Parameters

**src** (STRING): SnapMirror source volume.

## snapmirror.src.inopick.abort

### Severity

INFORMATIONAL

## Description

This message occurs when the SnapMirror® source appliance aborts a qtree transfer it was doing because the producer thread aborted, usually due to lack of memory.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**srcpath** (STRING): SnapMirror source path.  
**dstfiler** (STRING): SnapMirror destination appliance.  
**dstpath** (STRING): SnapMirror destination path.

## snapmirror.src.markedCorrupt

### Severity

ALERT

### Description

This event is issued when a SnapMirror is issued from a source volume or aggregate which is offline, restricted or does not exist.

### Corrective Action

(None).

### Syslog Message

Source %s %s is marked as corrupt, transfer not done.

### Parameters

**volaggr** (STRING): Whether volume or aggregate  
**src** (STRING): The SnapMirror source

## snapmirror.src.multipath.connErr

### Severity

ERROR

### Description

This event is issued when the SnapMirror source filer is unable to setup a multipath connection with the destination filer. This can happen in case of a network connection failure.

### Corrective Action

Check network connectivity. SnapMirror uses ports 10565 and 10566. Please make sure they are open in the firewall configuration.

### Syslog Message

SnapMirror is unable to set up a multipath/failover connection from %s to %s:%s, SnapMirror is resorting to a single TCP connection. Please ensure ports 10565 and 10566 are open in the firewall configuration.



## Parameters

**srcpath** (STRING): SnapMirror source  
**dstfiler** (STRING): SnapMirror destination filer  
**dstpath** (STRING): SnapMirror destination

## snapmirror.src.noNewData

### Severity

NOTICE

### Description

This generic event is issued when the SnapMirror source volume/aggregate itself is a SnapMirror destination and has no new data to send to its SnapMirror destination.

### Corrective Action

Please do a SnapMirror transfer to the source first.

### Syslog Message

SnapMirror source %s %s has no new data for %s:%s.

## Parameters

**volaggr** (STRING): Whether volume or aggregate  
**srcpath** (STRING): The SnapMirror source  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination

## snapmirror.src.notLicensed

### Severity

NOTICE

### Description

This event is issued when a SnapMirror source filer receives a request for transfer when SnapMirror is not even licensed on it.

### Corrective Action

Please license SnapMirror.

### Syslog Message

Command request while SnapMirror not licensed on this filer.

## Parameters

(None).

## snapmirror.src.notOn

### Severity

NOTICE

**Description**

This message occurs when a SnapMirror® source system receives a request for transfer when SnapMirror is off on the system.

**Corrective Action**

Issue the 'snapmirror on' command.

**Syslog Message**

Command request while SnapMirror is off on this filer.

**Parameters**

(None).

**snapmirror.src.notReady****Severity**

ERROR

**Description**

This event is issued when the SnapMirror source volume is not yet ready for a SnapMirror transfer.

**Corrective Action**

Please try again.

**Syslog Message**

Source %s %s not ready for SnapMirror transfer, please try again.

**Parameters**

**volaggr** (STRING): Whether volume or aggregate

**src** (STRING): The SnapMirror source

**snapmirror.src.offline****Severity**

ERROR

**Description**

This event is issued when the SnapMirror source volume or aggregate is taken offline.

**Corrective Action**

(None).

**Syslog Message**

Source %s %s requested offline.

**Parameters**

**volaggr** (STRING): Whether volume or aggregate

**src** (STRING): The SnapMirror source

## snapmirror.src.phasel

### Severity

NOTICE

### Description

This event is issued when the SnapMirror source doing a qtree transfer enters phase I of the transfer.

### Corrective Action

No corrective action required.

### Syslog Message

[ SnapMirror source transfer %s --> %s:%s ] starting phase I.

### Parameters

**srcpath** (STRING): The SnapMirror source path  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination path

## snapmirror.src.phasel1

### Severity

NOTICE

### Description

This event is issued when the SnapMirror source doing a qtree transfer enters phase II of the transfer.

### Corrective Action

No corrective action required.

### Syslog Message

[ SnapMirror source transfer %s --> %s:%s ] starting phase II.

### Parameters

**srcpath** (STRING): The SnapMirror source path  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination path

## snapmirror.src.requestDenied

### Severity

ERROR

### Description

This event is issued when the destination filer requesting a SnapMirror transfer from the source filer is denied the transfer request.

### Corrective Action

(None).

## Syslog Message

SnapMirror transfer request from %s to host %s at IP address %s denied: %s.

## Parameters

**srcpath** (STRING): The SnapMirror source, can be volume or qtree  
**hostname** (STRING): The hostname of the destination filer  
**ipaddress** (STRING): The IP address of the destination filer  
**error** (STRING): The reason why the request was denied.

## snapmirror.src.restart.chkptFailed

### Severity

NOTICE

### Description

This event is issued when we are not able to set a restart checkpoint on the qtree SnapMirror source.

### Corrective Action

(None).

## Syslog Message

SnapMirror transfer to %s:%s : restart checkpoint setup failed.

## Parameters

**dstfiler** (STRING): The SnapMirror destination filer  
**dstpath** (STRING): The SnapMirror destination qtree

## snapmirror.src.restart.noSnap

### Severity

ERROR

### Description

This event is issued when SnapMirror source is not able to find the latest snapshot needed for restart.

### Corrective Action

(None).

## Syslog Message

Could not find the latest snapshot needed for restart.

## Parameters

(None).

## snapmirror.src.resync.snapNotFound

### Severity

ERROR

## Description

This event is issued on the source when the resync base snapshot is not found on the source.

## Corrective Action

(None).

## Syslog Message

Could not find base snapshot to resync %s %s:%s to %s.

## Parameters

**volaggr** (STRING): Whether volume or aggregate  
**dstfiler** (STRING): SnapMirror destination filer  
**dstpath** (STRING): SnapMirror destination  
**srcpath** (STRING): SnapMirror source

## snapmirror.src.resync.toSelf

### Severity

ERROR

### Description

This event is issued when SnapMirror source detects and attempt to resync a volume to itself.

### Corrective Action

(None).

### Syslog Message

SnapMirror cannot resync %s %s to itself, operation not permitted.

### Parameters

**volaggr** (STRING): Where source volume or aggregate  
**srcpath** (STRING): SnapMirror source

## snapmirror.src.revertPending

### Severity

NOTICE

### Description

This message indicates the source is being reverted by WAFL and cannot currently be used for a snapmirror operation.

### Corrective Action

This message indicates the source is being reverted. After the revert completes the operation should be retried.

### Syslog Message

snapmirror copy: The source is being reverted by WAFL.

**Parameters**

(None).

**snapmirror.src.snapGotDeleted****Severity**

ERROR

**Description**

This event is issued when the base snapshot on the SnapMirror source got deleted.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror snapshot %s was deleted or renamed, transfer from %s to %s:%s aborted.

**Parameters**

**snapname** (STRING): The SnapMirror base snapshot

**srcpath** (STRING): The SnapMirror source, can be a volume or qtree

**dstfiler** (STRING): The SnapMirror destination filer

**dstpath** (STRING): The SnapMirror destination, can be a volume or qtree

**snapmirror.src.snapshotErr****Severity**

ERROR

**Description**

This event is issued when we are not able to obtain all the snapshot information, from the SnapMirror source volume, required for a SnapMirror transfer.

**Corrective Action**

(None).

**Syslog Message**

SnapMirror not able to obtain snapshot information from source %s %s

**Parameters**

**volaggr** (STRING): Whether volume or aggregate

**src** (STRING): The SnapMirror source volume

**snapmirror.src.sockErr****Severity**

ERROR

**Description**

This event is issued when SnapMirror on the source filer has problems opening a socket.

**Corrective Action**

(None).

**Syslog Message**

Cannot accept SnapMirror destination requests: %s.

**Parameters**

**error** (STRING): The error message

**snapmirror.src.sockOk****Severity**

NOTICE

**Description**

This event is issued when SnapMirror on the source filer had problems opening a socket, but is fine now.

**Corrective Action**

No corrective action required.

**Syslog Message**

SnapMirror now accepting destination requests.

**Parameters**

(None).

**snapmirror.src.sync.compressedVol****Severity**

ERROR

**Description**

This event is issued when a Synchronous SnapMirror source attempts to go into fully synchronous mode when compression is enabled on the source volume.

**Corrective Action**

SnapMirror currently does not support a compressed volume as a synchronous volume source. In order to synchronously mirror the volume, the compression feature must be disabled on the volume.

**Syslog Message**

Synchronous SnapMirror source volume %s has compression enabled, which is not supported.

**Parameters**

**src** (STRING): Name of the source volume.

**snapmirror.src.sync.flexSameFilerOrHA****Severity**

ERROR

## Description

This event is issued when a SnapMirror request for going synchronous is received for two flexible volumes that reside on the same filer or belong to the same HA pair.

## Corrective Action

Synchronous SnapMirror cannot operate on two flexible volumes on the same filer or on the same HA pair. Please move one of the volumes off the filer or HA pair.

## Syslog Message

Synchronous SnapMirror cannot run between two flexible volumes on the same filer or HA pair; this mode is not allowed from %s to %s.

## Parameters

**src** (STRING): The source volume name  
**dst** (STRING): The destination volume name

## snapmirror.src.sync.FvolSyncTooSmall

### Severity

ERROR

### Description

This event is issued when a source side SnapMirror request attempts to go into fully synchronous mode on a flexible volume and it is smaller than the supported size.

### Corrective Action

SnapMirror currently does not support a filer having a synchronous flexible volume source that is smaller than the minimum supported size. Increase the size of the flexible volume source and destination to at least the minimum supported size.

### Syslog Message

The flexible volume Synchronous SnapMirror source %s is %llu MB, which is smaller than the minimum supported size of %llu MB.

### Parameters

**src** (STRING): Name of the source flexible volume.  
**size** (LONGINT): Size of the source flexible volume (MB).  
**minsize** (LONGINT): Minimum supported volume size (MB) for flexible volume synchronous snapmirror transfers.

## snapmirror.src.sync.license

### Severity

NOTICE

### Description

This event is issued when a SnapMirror request for going synchronous is received when the feature is not licensed.

### Corrective Action

Please license synchronous SnapMirror on the source filer.



## Syslog Message

Synchronous SnapMirror is not licensed. Synchronous mode from %s to %s:%s is not allowed.

## Parameters

**srcvol** (STRING): The source volume name  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstvol** (STRING): The destination volume name

## snapmirror.src.sync.platform

### Severity

ERROR

### Description

This event is issued when a SnapMirror request for going synchronous is received when an inappropriate platform is used.

### Corrective Action

(None).

## Syslog Message

Synchronous SnapMirror is not supported on this platform. Synchronous mode from %s to %s:%s is not allowed.

## Parameters

**srcvol** (STRING): The source volume name  
**dstfiler** (STRING): The SnapMirror destination filer  
**dstvol** (STRING): The destination volume name

## snapmirror.src.sync.rootFullSync

### Severity

ERROR

### Description

This event is issued when a SnapMirror request for going fully synchronous is received for the root volume or aggregate.

### Corrective Action

(None).

## Syslog Message

Root %s %s not allowed as a fully sync source.

## Parameters

**volaggr** (STRING): Whether volume or aggregate  
**src** (STRING): The SnapMirror source

## snapmirror.src.sync.sanitization

**Severity**

NOTICE

**Description**

This event is issued when a SnapMirror request for going synchronous is received when SnapMirror features are only enabled due to a sanitization license.

**Corrective Action**

The sanitization feature only enables certain pieces of SnapMirror functionality. Please license synchronous SnapMirror in order to use synchronous mode.

**Syslog Message**

The sanitization feature does not allow synchronous SnapMirror to be run from volume %s to volume %s.

**Parameters**

**src** (STRING): The source volume name

**dst** (STRING): The destination volume name

**snapmirror.src.sync.tradSameFilerOrHA****Severity**

ERROR

**Description**

This event is issued when a SnapMirror request for going into fully synchronous mode is received for two traditional volumes that reside on the same filer or belong to the same HA pair.

**Corrective Action**

Fully synchronous SnapMirror does not operate on two traditional volumes on the same filer or on the same HA pair. The volumes can be put into semi-synchronous (i.e. cpsync) mode, or one of the volumes can be moved off the filer or HA pair.

**Syslog Message**

Fully synchronous SnapMirror cannot run between two traditional volumes on the same filer or HA pair; this mode is not allowed from %s to %s.

**Parameters**

**src** (STRING): The source volume name

**dst** (STRING): The destination volume name

**snapmirror.src.threadErr****Severity**

ERROR

**Description**

This event is issued when we are not able to allocate a thread to do SnapMirror source side operations.

**Corrective Action**

(None).

**Syslog Message**

Concurrent operation resource limit reached, SnapMirror source delayed.

**Parameters**

(None).

**snapmirror.src.versionMismatch****Severity**

ERROR

**Description**

This event is issued when there is SnapMirror version mismatch between the source and destination.

**Corrective Action**

(None).

**Syslog Message**

Received SnapMirror request with incompatible version.

**Parameters**

(None).

**snapmirror.src.vfiler.migrating****Severity**

ERROR

**Description**

This message occurs when a SnapMirror® transfer involves a source volume that is a member of a migrating vFiler(tm) unit. In such a case, the SnapMirror transfer aborts.

**Corrective Action**

Wait for the vFiler migration to be completed, then retry the transfer.

**Syslog Message**

SnapMirror source volume %s is member of a migrating vFiler unit.

**Parameters**

**src** (STRING): SnapMirror source volume

**snapmirror.src.vfilerAccess****Severity**

ERROR

**Description**

This event occurs when a SnapMirror source running on a vfiler tries to access a volume that it does not own.

## Corrective Action

For a non-default vfiler to act as a SnapMirror source it must own the target volume. Do one of the following: a) use the 'vfiler move' command to make the given vfiler the owner of the volume, b) connect to the default vfiler that has access to all volumes, or c) connect to the vfiler that currently owns the specified SnapMirror source volume. See the na\_vfiler man page and MultiStore® documentation for details.

## Syslog Message

SnapMirror : vfiler %s tried to use volume %s as a SnapMirror source volume, but the vfiler is not the owner of that volume.

## Parameters

**vfiler** (STRING): The current vfiler's name  
**src** (STRING): The SnapMirror source volume

# snapmirror.status events

## snapmirror.status.dstUpdateSnapErr

### Severity

ERROR

### Description

This event is generated when Data ONTAP cannot update the latest base snapshot for the replication destination. Insufficient disk space on the root volume is the most common reason for this failure.

### Corrective Action

Check if the root volume is out of disk space.

### Syslog Message

Data ONTAP could not update the registry with latest snapshot for the destination %s.

### Parameters

**dstPath** (STRING): The destination path

## snapmirror.status.illegalSrcPath

### Severity

ERROR

### Description

This message occurs on a replication source appliance, when Data ONTAP encounters invalid source path specification while collecting replication statistics.

### Corrective Action

Check snapmirror.conf file for any typing mistakes. Typically path names are specified as "filer:volume" or "filer:/vol/volume/qtrees". Check snapmirror.conf man page for discussion of path specifications.

### Syslog Message

Illegal source path: %s.

## Parameters

**srcPath** (STRING): The Source path.

## snapmirror.status.noBaseSnapshot

### Severity

ERROR

### Description

This message occurs on replication destination appliance when Data ONTAP can not find the base snapshot used by a replication relation. Data ONTAP may run into this situation while updating replication statistics at the end of a successful transfer.

### Corrective Action

This is most likely a temporary error. Please try the operation again. Check status of the destination volume.

### Syslog Message

Can not find base snapshot for %s.

## Parameters

**dstPath** (STRING): The destination path.

## snapmirror.status.updateStatusErr

### Severity

ERROR

### Description

This event is generated when Data ONTAP cannot update the ONTAP system registry with upgraded snapmirror status information. Insufficient disk space on the root volume is the most common reason for this failure.

### Corrective Action

Check if the root volume is out of disk space by issuing the 'df' command from the appliance CLI. If the root volume is full, either free up space or add more disks on the volume.

### Syslog Message

Data ONTAP could not update the ONTAP system registry with upgraded snapmirror status keys.

## Parameters

(None).

## snapmirror.tradtoflexerr events

### snapmirror.tradToFlexErr

### Severity

ERROR

**Description**

This message occurs when the user tries to copy a traditional volume into a flexible volume using SnapMirror®.

**Corrective Action**

Retry the operation with a traditional volume as the destination.

**Syslog Message**

snapmirror : A traditional volume cannot be copied into a flexible volume.

**Parameters**

(None).

## snapmirror.unsupportedchecksumtype events

### snapmirror.unsupportedChecksumType

**Severity**

NOTICE

**Description**

This message occurs when checksum type `crc32c_header_only` is configured for a SnapMirror® relationship and the source for the relationship is not compatible with this checksum type.

**Corrective Action**

(None).

**Syslog Message**

%s: destination transfer from %s:%s to %s : %s

**Parameters**

**rtype** (STRING): Replication type (SnapMirror).  
**srcfiler** (STRING): Replication source system.  
**srcpath** (STRING): Replication source volume.  
**dstpath** (STRING): Replication destination volume.  
**message** (STRING): Specific message.

# snapmirror\_copy events

## snapmirror\_copy.updatetimeincreased events

### snapmirror\_copy.UpdateTimeIncreased

#### Severity

NOTICE

#### Description

This message occurs when the update time taken by the latest transfer is 5 minutes more than the previous three transfer times.

#### Corrective Action

(None).

#### Syslog Message

Update time for volume %s has increased.

#### Parameters

**dstpath** (STRING): Destination volume.

# snaprestore events

## snaprestore.fail events

### snaprestore.fail.volume.size

#### Severity

ALERT

#### Description

This message occurs when a SnapRestore® operation on a volume fails due to insufficient space in the target volume to hold the data present in the Snapshot(tm) copy.

#### Corrective Action

Resize the target volume to the recommended minimum size.

#### Syslog Message

SnapRestore operation failed: volume %s is too small to hold the data present in the Snapshot copy. Minimum size of the target volume required is %s.

#### Parameters

**volume** (STRING): Volume name.

**volsize** (STRING): Recommended minimum size of the target volume.

## snaprestore.log events

### snaprestore.log.overalloc

#### Severity

EMERGENCY

#### Description

This message occurs during a SnapRestore™ operation, of a volume when the logical used space in the volume is greater than the total size of the volume and the Logical Space View enforcement feature was turned ON on the volume after creating the Snapshot™ copy. This might cause client modify operations to fail on the volume.

#### Corrective Action

Create space by increasing the volume size, or by deleting data or deleting Snapshot® copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot® copies, use the "volume snapshot delete" command.

#### Syslog Message

During SnapRestore operation, volume %s in aggregate %s was logically overallocated during snap restore operation. It may impact client modify operations.

#### Parameters

**volume** (STRING): Volume name.

**aggr** (STRING): Name of the containing aggregate.



# snaptags events

## snaptags.file events

### snaptags.file.inconsist

**Severity**

NOTICE

**Description**

This message occurs when the system detects an inconsistency in the header of a Snapshot(tm) copy tag's metafile of a volume. The Snapshot copy tags for all of the Snapshot copies on the volume are considered lost. The next Snapshot copy tag update starts from a clean metafile.

**Corrective Action**

(None).

**Syslog Message**

An inconsistency was detected in the Snapshot copy tag's metafile on volume '%s'. The Snapshot copy tags for all of the Snapshot copies on the volume are lost.

**Parameters**

**vol** (STRING): Name of the volume.

## snaptags.snap events

### snaptags.snap.inconsist

**Severity**

NOTICE

**Description**

This message occurs when the system detects inconsistency of Snapshot(tm) copy-specific information in the Snapshot copy tag's metafile on a volume. The tags for the Snapshot copy are considered lost. The next tag update for the Snapshot copy starts without any existing tag.

**Corrective Action**

(None).

**Syslog Message**

An inconsistency of Snapshot copy tags was detected for Snapshot copy ID '%d' on volume '%s'. The tags for the Snapshot copy are lost.

**Parameters**

**snapid** (INT): Physical Snapshot copy ID on the volume.

**vol** (STRING): Name of the volume.

# snapvault events

## snapvault.dst events

### snapvault.dst.bkpSnapCreate

#### Severity

INFORMATIONAL

#### Description

This message occurs when SnapVault® creates a backup Snapshot(tm) copy for a given backup schedule.

#### Corrective Action

(None).

#### Syslog Message

A backup Snapshot copy was created on volume %s by backup schedule %s (%d of %d allocated Snapshot copies consumed).

#### Parameters

**vol\_name** (STRING): Name of the SnapVault secondary volume on which the backup Snapshot copy is created.

**target\_name** (STRING): Base name of the backup schedule for which the given backup Snapshot copy is created.

**snap\_cnt** (INT): Total number of backup Snapshot copies created for the given backup schedule.

**alloc\_cnt** (INT): Total number of backup Snapshot copies allocated for the given backup schedule.

### snapvault.dst.bliDataRepaired

#### Severity

NOTICE

#### Description

This message is issued when SnapVault may have repaired data that had previously been made incorrect/incomplete by Bug #137685 which has now been fixed. Under some conditions this message may be issued when, in fact, the data was not previously incorrect or incomplete. Under some conditions, the SnapVault code cannot determine whether the previous version of the data was correct or not, but it does know that after the current transfer completes that the data is correct. This message will only be issued if the hidden snapvault option, repair\_bug137685\_stats is enabled.

#### Corrective Action

(None).

#### Syslog Message

Data ONTAP changed data in %d data chunk(s) in %s which may have been missing or incorrect.

#### Parameters

**chunks\_repaired** (INT): The number of of chunks which may have been repaired. Up to 2MB of data may be sent in a chunk in the QSM/SnapVault data format. This count is incremented for each chunk in which one or more 4K blocks may have been repaired.

**dst\_path** (STRING): Destination qtree path.

## snapvault.dst.bliWarning

### Severity

ERROR

### Description

This event is issued when SnapVault has detected that OSSV secondary data may be incomplete or incorrect.

### Corrective Action

See the public reports for Bug #137685 and Bug #140930 for corrective action.

### Syslog Message

Data ONTAP has detected a condition in which OSSV secondary data for %s may be incomplete or incorrect. See the public reports for Bug #137685 and Bug #140930 for corrective action.

### Parameters

**dst\_path** (STRING): Destination qtree path.

## snapvault.dst.lowSnapWarn

### Severity

ERROR

### Description

This message occurs when the remaining number of Snapshot(tm) copies allocated for a given SnapVault® backup schedule on the given volume reaches or falls below the preset warning level. This event occurs on every failed attempt by SnapVault to create a backup Snapshot copy on the given backup schedule, until you resolve the condition.

### Corrective Action

Consider deleting some backup Snapshot copies from the given backup target to create some free slots. For more information and other solutions, see the Online Backup Recovery Guide.

### Syslog Message

Remaining number of Snapshot copies on the volume %s for the backup schedule %s is below the set warning level (%d of allocated %d).

### Parameters

**vol\_name** (STRING): Name of the SnapVault secondary volume on which the backup Snapshot copy is created.

**target\_name** (STRING): Base name of the backup schedule for which the given backup Snapshot copy is created.

**snap\_cnt** (INT): Total number of backup Snapshot copies created for the given backup schedule.

**alloc\_cnt** (INT): Total number of backup Snapshot copies allocated to the given backup schedule.

## snapvault.dst.opNotComplete

### Severity

NOTICE

## Description

This message occurs when SnapVault® cannot complete a transfer normally, due to resource unavailability or SnapVault being off.

## Corrective Action

None needed. SnapVault periodically attempts to complete this transfer.

## Syslog Message

(None).

## Parameters

**dst\_path** (STRING): Destination qtree path.

**snapvault\_on** (INT): SnapVault enabled? - on(1) or off(0)

## snapvault.dst.updateDelayed

### Severity

INFORMATIONAL

### Description

This message occurs when an update request is received on a destination qtree that is being used as the source for another transfer.

### Corrective Action

(None).

### Syslog Message

%s: Update from %s:%s to %s:%s is delayed because another update, %s:%s to %s:%s, is already in progress.

### Parameters

**type** (STRING): SnapVault® transfer type.

**srcfiler1** (STRING): SnapVault source system.

**srcpath1** (STRING): SnapVault source qtree.

**dstfiler1** (STRING): SnapVault destination system for first transfer (in progress). Source for second transfer (delayed).

**dstpath1** (STRING): SnapVault destination qtree for first transfer. Source for second transfer.

**srcfiler2** (STRING): SnapVault destination system for first transfer. Source for second transfer.

**srcpath2** (STRING): SnapVault destination qtree for first transfer. Source for second transfer.

**dstfiler2** (STRING): Destination system for second transfer.

**dstpath2** (STRING): Destination qtree for second transfer.

## snapvault.qtree events

### snapvault.qtree.initiate

#### Severity

NOTICE

## Description

This message occurs when SnapVault® cannot start a volume-wide Snapshot(tm) target to complete a manually invoked baseline transfer or resynchronizing transfer. As a result, the changes to the qtree in that operation will not yet be made visible. SnapVault attempts to repair this condition.

## Corrective Action

(None).

## Syslog Message

Could not coalesce %s: %s.

## Parameters

**qtpath** (STRING): Path of the qtree in the operation.

**reason** (STRING): Reason why SnapVault could not start the Snapshot target.

## snapvault.qtree.notCoalesced

### Severity

ERROR

### Description

This message occurs when SnapVault® fails to initiate a qtree into a SnapVault coalescing cycle.

### Corrective Action

If the error is transient, the next SnapVault coalescing cycle tries to initiate the qtree into the coalescing cycle. Otherwise, see the error message for details.

### Syslog Message

Could not initiate %s into SnapVault coalescing cycle: %s.

### Parameters

**qtpath** (STRING): Path of the qtree in the operation.

**errstr** (STRING): Error string.

## snapvault.qtree.upgradeStatus

### Severity

NOTICE

### Description

During initialization, SnapVault will attempt to update the qtree metadata flag if necessary. This message is generated if the qtree has been converted but an error occurs in updating the status entry.

### Corrective Action

(None).

### Syslog Message

Could not update the status entry for %s. Entry will be updated during the next transfer.

## Parameters

**srcpath** (STRING): The pathname for the qtree.

# snapvault.reg events

## snapvault.reg.unknownOpt

### Severity

NOTICE

### Description

This message occurs when you downgrade from a kernel with newer version to one with an older version and the system attempts to read a registry entry that is supported only in the newer version of the kernel.

### Corrective Action

(None).

### Syslog Message

Unsupported registry entry: (%s).

## Parameters

**regopt** (STRING): Unrecognized registry entry for the older version of the kernel.

# snapvault.src events

## snapvault.src.notEnabled

### Severity

ERROR

### Description

This message occurs when a SnapVault® source appliance receives a transfer request when the service is not enabled. The license needed depends on the transfer type. For SnapVault restore transfers, the appliance must have a secondary license. For other transfers, the appliance must have a primary license.

### Corrective Action

Enable the service by issuing the command 'options snapvault.enable on'. You might need to license the service before you enable it.

### Syslog Message

Command request (%d) from host %s while SnapVault is not enabled.

## Parameters

**cmd** (INT): Protocol command number for the type of transfer being requested.

**dst\_ip** (STRING): IP address of the destination issuing the transfer request.

## snapvault.src.volume.migrating

**Severity**

ERROR

**Description**

This message occurs when a SnapVault® Snapshot(tm) target is initiated on a SnapVault primary volume while the volume is being migrated. The target might have been scheduled or manually invoked (through the 'snapvault snap create' command or programmatic interfaces).

**Corrective Action**

Wait for the volume migration to be completed.

**Syslog Message**

SnapVault primary volume %s is currently being migrated.

**Parameters**

**volname** (STRING): Name of the SnapVault primary volume, on which the target is invoked.

## snapvault.sys events

### snapvault.sys.internal

**Severity**

ERROR

**Description**

This message occurs when SnapVault® encounters an internal error.

**Corrective Action**

Turn off SnapVault and restart it.

**Syslog Message**

An internal error occurred in SnapVault@: %s.

**Parameters**

**reason** (STRING): Reason for the internal error.

### snapvault.sys.upgrade

**Severity**

ERROR

**Description**

During initialization, SnapVault will attempt to update all the qtree metadata flags if necessary. This message is generated if there is insufficient memory to perform this operation.

**Corrective Action**

Convert the qtrees manually or contact NetApp technical support to help.

**Syslog Message**

Low on memory. Can't automatically convert snapvault qtrees right now, so backups will fail. Convert the qtrees manually or contact NetApp technical support to help.

## Parameters

(None).

## snapvault.sys.vfilerOn

### Severity

INFORMATIONAL

### Description

This message occurs when SnapVault® is enabled on a vfiler during shutdown, where it is expected to be disabled.

### Corrective Action

None needed. This is for information only.

### Syslog Message

(None).

### Parameters

**vfiler** (STRING): Name of the vfiler.

## snapvault.tgt events

### snapvault.tgt.dupCoalesced

#### Severity

NOTICE

#### Description

This message occurs when SnapVault® detects more than one coalesced Snapshot(tm) copy of a specific type in a SnapVault secondary volume. SnapVault attempts to repair this condition.

#### Corrective Action

(None).

#### Syslog Message

Detected multiple coalesced %s Snapshot copies in volume %s.

#### Parameters

**snaptype** (STRING): Type of coalesced Snapshot copy that has duplicates.

**volume** (STRING): Name of the SnapVault secondary volume on which the target was invoked.

### snapvault.tgt.failure

#### Severity

ERROR

#### Description

This message occurs when a SnapVault® Snapshot(tm) target fails, either because it could not start or because it encountered an error while processing. The target could have been scheduled or manually



invoked (through the 'snapvault snap create' command or programmatic interfaces).

### Corrective Action

The corrective action depends on the cited reason regarding why the Snapshot target could not start or why it failed. After you correct the situation, invoke the Snapshot target again, or wait for the next scheduled invocation. If the volume could not be found, make sure that it is online, check the spelling of the volume name, and correct the SnapVault Snapshot schedule if it is misspelled. If the target could not start due to a temporary lack of resources (such as threads or memory), no corrective action might be necessary, because SnapVault periodically retries the operation. If the problem persists, you might need to schedule SnapVault and/or other services to run less frequently. If the target could not create a Snapshot copy, or could not create or write either data or metadata to the volume, check whether the volume has any available space, inodes, and Snapshot copies. If any of these have been exhausted, free some up for SnapVault use. If the target didn't run because it is not a create target, install a SnapVault secondary license or unconfigure the target. If the target didn't run because some other service is actively using the volume and would interfere with SnapVault, either stop that service or wait for it to be completed.

### Syslog Message

Could not create Snapshot target "%s" on volume %s: %s.

### Parameters

**snapshotname** (STRING): Base name of the Snapshot copy that the target was meant to create.

**volume** (STRING): Name of the SnapVault secondary volume on which the target was invoked.

**reason** (STRING): Reason why the target could not start, or the specific error encountered during processing.

## snapvault.tgt.NBUfailure

### Severity

ERROR

### Description

This event occurs when a snapshot creation on the destination volume used to back up primary qtrees using SnapVault(tm) managed by VERITAS® NetBackup® fails. Because the snapshot creation failure also causes the NetBackup backup request to fail, this is considered an error. This failure occurs when the number of snapshots in the volume exceeds the maximum limit, there is insufficient disk space in the destination volume, or the destination volume is offline or restricted.

### Corrective Action

Use the 'snap list volname' command to list the snapshots in the volume and to see whether there are close to 255 snapshots (the maximum). Use the 'df -k' command to check the amount of disk space available in the volume. If either of these conditions is the cause of the failure, contact the NetBackup Administrator and ask for some old backups stored on this volume to be deleted. This results in both the number of snapshots being reduced and disk space getting freed. Also check the status of the volume by using the 'vol status' command to confirm that the volume is offline or restricted. If it is, use the 'vol online volname' command to bring the volume online.

### Syslog Message

SnapVault snapshot creation for volume (%s) failed.

### Parameters

**volume** (STRING): The name of the destination volume for which the snapshot creation failed.

## snapvault.tgt.qtreeignore

### Severity

ERROR

### Description

This message occurs when SnapVault® intentionally ignores a configuration entry for a secondary qtree. After ignoring it once, SnapVault continues ignoring the qtree (silently) until you update the qtree manually, until you restart the SnapVault service, or until the next boot.

### Corrective Action

The corrective action depends on the cited reason regarding why the qtree configuration was ignored. Frequently, the qtree simply doesn't exist, in which case the configuration entry is stale and should be removed (through the 'snapvault stop' command). If the qtree was ignored due to some transient reason, you need to manually update the qtree, or disabled and reenables SnapVault to begin operating on the qtree again.

### Syslog Message

Ignoring configuration entry for %s in volume %s: %s.

### Parameters

**qtpath** (STRING): Path to the qtree whose configuration entry was ignored.

**volume** (STRING): Name of the SnapVault secondary volume with the ignored configuration entry.

**reason** (STRING): Reason why the qtree configuration was ignored.

## snapvault.tgt.unstableSnap

### Severity

ERROR

### Description

This message occurs when SnapVault® detects a coalesced Snapshot(tm) copy with transitioning qtree replicas. If this event occurs, it suggests that archive Snapshot copies might have captured some SnapVault secondary qtrees in a transitioning state, which you can verify by running the 'snap list -q' command. SnapVault attempts to recover from this situation, but that only affects new archive Snapshot copies created after the event, and not any existing Snapshot copies.

### Corrective Action

(None).

### Syslog Message

Recovery detected %s Snapshot copy with transitioning qtrees in volume %s.

### Parameters

**snaptype** (STRING): Type of coalesced Snapshot copy that contains transitioning SnapVault qtree replicas.

**volume** (STRING): Name of the SnapVault secondary volume on which the target was invoked.

# snmp events

## snmp.authentication events

### snmp.authentication.failure

#### Severity

ERROR

#### Description

This message occurs when there is an authentication failure during an SNMP query. This trap is generated for all SNMP versions.

#### Corrective Action

In case of SNMPv1/v2c check whether the SNMP community string is configured. In case of SNMPv3, check whether SNMPv3 authentication credentials are correctly configured. The SNMPv3 credentials include username, password, authentication protocol, authentication protocol password, privacy protocol and privacy protocol password.

#### Syslog Message

Authentication failure for SNMP query from port: %d.

#### Parameters

**interface** (INT): Port on which the SNMP authentication failed.

## snmp.coldstart events

### snmp.coldstart.trap

#### Severity

INFORMATIONAL

#### Description

This message occurs when the SNMP daemon is reinitializing itself with a coldStart. A coldStart(0) trap signifies that configuration changes are made to the agent or the protocol entity implementation during the reinitialization.

#### Corrective Action

(None).

#### Syslog Message

SNMP daemon was reinitialized with configuration changes.

#### Parameters

(None).

## snmp.fips events

## snmp.fips.objs.del.failed

### Severity

ERROR

### Description

This message occurs when the system fails to delete SNMP traphosts and users that are not FIPS compliant. The system tries to automatically delete them when FIPS is configured in the cluster. The SNMP traphosts and users that are left become inoperable.

### Corrective Action

Delete remaining noncompliant SNMP traphosts first by using the "system snmp traphost delete" command. Then delete the remaining noncompliant SNMP users by using the "security login delete" command. The following SNMP users and traphosts are not FIPS compliant: 1. SNMPv1 users, SNMPv2c users, and SNMPv3 users having "none" or "MD5" as the authentication method or having "none" or "DES" as the encryption protocol. 2. SNMPv1 traphosts and SNMPv3 traphosts configured with a user that is not FIPS compliant.

### Syslog Message

Failed to delete SNMP traphosts and users that are not FIPS compliant.

### Parameters

(None).

## snmp.fips.support

### Severity

INFORMATIONAL

### Description

This message occurs when Data ONTAP is operating in FIPS compliant mode but SNMPv3 is configured with non-compliant/weaker ciphers or hash algorithms. Hence SNMPv3 is disabled.

### Corrective Action

(None).

### Syslog Message

Data ONTAP is operating in FIPS compliant mode and SNMPv3 is configured with non-compliant ciphers or hash algorithms. Hence SNMPv3 is disabled.

### Parameters

(None).

## snmp.link events

### snmp.link.down

### Severity

INFORMATIONAL

## Description

This message occurs when a network interface is down.

## Corrective Action

If this message is not followed by a LinkUp message, the cluster administrator should manually check the link state with the "network port show" command, verifying that the port is administratively up. If the link state is still down and the port is not intentionally disabled, there might be a physical network problem: 1. Verify the switch has the corresponding port administratively and operationally up. 2. Verify that the network cable between the storage controller port and the corresponding switch is connected to the correct port and is not damaged. 3. Verify the integrity of the network cable with a cable tester, or swap in another cable known to be in good condition. 4. For further information or assistance contact NetApp technical support

## Syslog Message

Interface %d is down.

## Parameters

**interface** (INT): Number of the interface.

## snmp.link.up

### Severity

INFORMATIONAL

## Description

This message occurs when a network interface is up.

## Corrective Action

(None).

## Syslog Message

Interface %d is up

## Parameters

**interface** (INT): The number of the interface

## snmp.raid events

### snmp.raid.iterator.failure

### Severity

ERROR

## Description

This event is generated when a RAID group iterator was not able to be created.

## Corrective Action

Check for other resource related issues.

## Syslog Message

SNMP failed to create RAID group iterator

## Parameters

(None).

# snmp.server events

## snmp.server.busy

### Severity

ERROR

### Description

This message occurs when the SNMP daemon is busy serving a high rate of incoming SNMP queries. This might be caused by multiple SNMP clients querying the cluster or a malicious client attempting a denial-of-service attack. SNMP clients might experience delays or timeouts.

### Corrective Action

Run "tcpdump" on cluster to determine: If multiple SNMP clients querying the cluster, then reduce the rate of SNMP queries to the cluster or stop them entirely for a period of time (~15 minutes). If a malicious SNMP client is trying a denial-of-service attack, then update the firewall policy to block the client from querying the cluster by using the "firewall policy modify -policy mgmt -service snmp -action deny -ip-list IP\_address" command.

### Syslog Message

SNMP daemon is busy serving a high rate of incoming SNMP queries.

## Parameters

(None).

# snmp.snmpv3 events

## snmp.snmpv3.enable

### Severity

ERROR

### Description

This message occurs when SNMPv3 is disabled due to a cluster in FIPS mode being upgraded to a version of ONTAP® software supporting the SNMPv3 FIPS feature. FIPS compliance adds more stringent limitations on SNMP users and SNMP traphosts, so SNMPv3 access has been disabled to allow the current SNMP users and SNMP traphosts to be reevaluated. If SNMPv3 functionality is required, you can reenableView it; if cluster is still in FIPS mode at that time, any existing SNMP users and SNMP traphosts that are noncompliant to FIPS will be deleted. Until SNMPv3 is reenableViewed, SNMPv3 users cannot be created, SNMPv3 traphosts cannot be added and SNMPv3 packets will not be processed.

### Corrective Action

Enable SNMPv3 in the cluster by using the 'system snmp enable-snmpv3' command.

### Syslog Message

SNMPv3 is disabled due to a cluster in FIPS mode being upgraded to a version of ONTAP® software supporting the SNMPv3 FIPS feature.

## Parameters

(None).

# snmp.traphost events

## snmp.traphost.resolve.failed

### Severity

ERROR

### Description

The snmp daemon could not resolve (find an IP address via DNS for) a traphost. The appliance will try again in an hour to resolve the traphost.

### Corrective Action

Make sure that the network, including the domain server and traphost, are connected, running, and configured correctly. If you fix this problem, wait until the next retry or run 'snmp init 1' to force a resolve. If you don't fix this problem, you may want to remove this traphost from the list. Attempting to resolve traphosts may take time because of network timeouts.

### Syslog Message

snmp: cannot send traps to '%s' because it could not be resolved via DNS. Retries occur hourly.

## Parameters

**host** (STRING): The host that could not be resolved via a DNS lookup.

# snmp.warmstart events

## snmp.warmstart.trap

### Severity

INFORMATIONAL

### Description

This message occurs when the SNMP daemon is reinitializing itself with a warmStart. A warmStart(1) trap signifies that no configuration changes are made to the agent or the protocol entity implementation during the reinitialization.

### Corrective Action

(None).

### Syslog Message

SNMP daemon was reinitialized with no configuration changes.

## Parameters

(None).

# socket events

## socket.limit events

### socket.limit.exceeded

#### Severity

NOTICE

#### Description

This message occurs when the operating system is unable to allocate a new socket because that would exceed the maximum open sockets.

#### Corrective Action

(None).

#### Syslog Message

The operating system was unable to allocate a socket. The number of open sockets is %d and the limit is %d.

#### Parameters

**num\_open\_sockets** (INT): The number of open sockets.

**max\_sockets** (INT): The maximum allowed number of open sockets.

## socket.sndbuflimit events

### socket.sndbuflimit.exceeded

#### Severity

ERROR

#### Description

This message occurs when an overflow is detected in the TCP socket send buffer. The TCP socket send buffer can hold a maximum of INT\_MAX (2147483647) bytes.

#### Corrective Action

Excessive unacknowledged data in the socket send buffer is commonly due to an overloaded client that is extremely slow to send TCP acknowledgments or completely stops doing so. Examine the client and take actions to address the client behavior.

#### Syslog Message

The system (%s:%d) could not send data to client %s due to an overflow in the TCP socket send buffer. The buffer can hold a maximum of %d bytes.

#### Parameters

**server\_address** (STRING): IP address of the server.

**port** (INT): Local port bound to the socket.

**client\_address** (STRING): IP address of the client.

**socket\_buffer\_limit** (INT): Maximum allowed number of bytes in socket buffer.



# sp events

## sp.auth events

### sp.auth.loginDenied

**Severity**

NOTICE

**Description**

This message occurs when a login is attempted on the node's Service Processor (SP) or baseboard management controller (BMC) and the login is denied.

**Corrective Action**

(None).

**Syslog Message**

BMC or SP login authentication failure for user "%s" from %s.

**Parameters**

**user** (STRING): User the login was attempted with.

**fromhost** (STRING): Where the login was attempted from.

## sp.authlog events

### sp.authlog.failure

**Severity**

ERROR

**Description**

This message occurs when the Service Processor (SP) or baseboard management controller (BMC) fails to log an authentication event in the system event log (SEL).

**Corrective Action**

1. Use the "system service-processor reboot-sp" command to reboot the SP or BMC. 2. If the failure persists, contact NetApp technical support

**Syslog Message**

BMC or SP failed to log an authentication event in the SEL.

**Parameters**

(None).

## sp.autoconfig events

## sp.autoconfig.out.of.address

### Severity

ERROR

### Description

This message occurs when the system cannot automatically obtain IP configuration because the subnet associated with the Service Processor does not have any available IP addresses.

### Corrective Action

Expand the subnet range to include more IP addresses with command "network subnet add-ranges". If failures persist, contact NetApp technical support.

### Syslog Message

(None).

### Parameters

**subnet** (STRING): Name of the subnet associated with the Service Processor.

## sp.autoupdate events

### sp.autoupdate.toggle

### Severity

INFORMATIONAL

### Description

This message occurs when the Service Processor (SP) AutoUpdate option is enabled or disabled.

### Corrective Action

(None).

### Syslog Message

%s.

### Parameters

**reason** (STRING): SP AutoUpdate feature has changed state.

## sp.device events

### sp.device.fw.mismatch

### Severity

ERROR

### Description

This message occurs when the service processor (SP) detects a firmware version mismatch among the devices managed by the SP.

### **Corrective Action**

Contact NetApp technical support to update the firmware.

### **Syslog Message**

Device firmware version mismatch detected: %s

### **Parameters**

**mismatch** (STRING): Description of the device names, current device firmware versions, and expected versions.

## **sp.dhcp events**

### **sp.dhcp.failure**

#### **Severity**

ERROR

#### **Description**

This message occurs when the Service Processor (SP) cannot obtain network configuration data through DHCP. The SP cannot be reached from the network.

#### **Corrective Action**

Verify that DHCP services are available, or configure the SP with a static address. If failures persist, contact NetApp technical support.

#### **Syslog Message**

Service Processor cannot obtain network configuration using DHCP.

#### **Parameters**

(None).

## **sp.dup events**

### **sp.dup.ip**

#### **Severity**

ERROR

#### **Description**

This message occurs when another system has been detected to be using the same IP address as the one configured on the Service Processor (SP).

#### **Corrective Action**

If the network in question has been configured to use this address as a static address, change the configured IP address on at least one system to ensure that only one system is using this address. If the network is configured to use DHCP, ensure that this address is not a static IP address that falls within the range being used by the DHCP server.

#### **Syslog Message**

SP IP address %s in use by another system.

## Parameters

**ipaddr** (STRING): IP address.

# sp.eth events

## sp.eth.switch.upd.failed

### Severity

ERROR

### Description

This message occurs when Service Processor (SP) or Baseboard Management Controller (BMC) fails to auto-update the internal Ethernet switch.

### Corrective Action

1. Perform a manual Ethernet Switch update from SP or BMC CLI.
2. If the failures persist, contact NetApp technical support.

### Syslog Message

%s Ethernet switch auto-update failed.

### Parameters

**sp\_type** (STRING): SP type

## sp.eth.switch.upd.start

### Severity

NOTICE

### Description

This message occurs when the Service Processor (SP) or Baseboard Management Controller (BMC) starts an internal Ethernet switch auto-update. During the Ethernet switch update, SP or BMC and e0M network connections might be temporarily affected.

### Corrective Action

(None).

### Syslog Message

%s is starting Ethernet switch auto-update %s.

### Parameters

**sp\_type** (STRING): SP type

**switch\_ver** (STRING): Current and bundled switch version.

## sp.eth.switch.upd.success

### Severity

NOTICE

## Description

This message occurs when the Service Processor (SP) or Baseboard Management Controller (BMC) completes an internal Ethernet switch auto-update.

## Corrective Action

(None).

## Syslog Message

%s Ethernet switch version updated %s.

## Parameters

**sp\_type** (STRING): SP type

**switch\_ver** (STRING): Previous and updated switch version.

# sp.firmware events

## sp.firmware.corruption

### Severity

ALERT

### Description

This message occurs when neither of the firmware images of the node's Service Processor (SP) or Baseboard Management Controller (BMC) on the primary flash chip can be booted. The SP or BMC is now running from the image on the spare flash storage. The SP or BMC spare image does not have all the features of the primary image, and it does not support updating the primary flash images.

### Corrective Action

Contact NetApp technical support to recover the primary flash chip.

### Syslog Message

"%s" is running from the spare recovery flash storage.

### Parameters

**sp\_type** (STRING): Service processor type (SP or BMC).

## sp.firmware.incompatible

### Severity

ALERT

### Description

This message occurs when the firmware on the Service Processor (SP) is not compatible with this version of Data ONTAP®.

### Corrective Action

Update the SP firmware to a version supported by this version of Data ONTAP. Contact NetApp technical support to obtain SP firmware files and update instructions, and for further assistance if the error persists after SP firmware is updated. Additional support might be available from the vendor's support site or knowledge base.

## Syslog Message

The SP firmware version %s is not compatible with this version of Data ONTAP.

## Parameters

**current\_sp\_fw\_version** (STRING): Current firmware version on the SP.

## sp.firmware.upgrade.reqd

### Severity

ALERT

### Description

This message occurs when the Service Processor (SP) firmware version and the Data ONTAP® software version are incompatible and cannot communicate correctly about a particular capability.

### Corrective Action

Update the firmware version of the SP to the version recommended for your version of Data ONTAP. Contact NetApp technical support to get a copy of the appropriate SP firmware and SP firmware update instructions. For more information about downloading and installing updated firmware, check your support provider's support web site or knowledgebase. If you need further assistance, contact NetApp technical support. After you update the firmware, this message should no longer occur. If the message occurs again, contact NetApp technical support and explain that you already updated the firmware to the recommended version.

## Syslog Message

The SP firmware %s is incompatible with Data ONTAP for %s.

## Parameters

**current\_sp\_fw\_version** (STRING): Current firmware version on the SP.

**capability** (STRING): SP capability needed by Data ONTAP.

## sp.firmware.version.unsupported

### Severity

ALERT

### Description

This message occurs when the firmware on the Service Processor (SP) is not an official release and must be updated.

### Corrective Action

Contact NetApp technical support to get a copy of the recommended version of SP firmware and SP firmware update instructions. For more information about downloading and installing updated firmware, check your support provider's web site or knowledgebase. If you need further assistance, contact NetApp technical support. After the SP is running the new firmware, this message should no longer occur. If the message occurs again, contact NetApp technical support and explain that you already updated the firmware to the recommended version.

## Syslog Message

The SP firmware version %s is not an official release.

## Parameters

**current\_sp\_fw\_version** (STRING): Current firmware version on the SP.

# sp.heartbeat events

## sp.heartbeat.resumed

### Severity

INFORMATIONAL

### Description

This event is generated when the appliance detects resumption of Service Processor (SP) heartbeat notifications indicating that the Service Processor is now available. The earlier issue indicated by the sp.heartbeat.stopped event has been resolved.

### Corrective Action

(None).

### Syslog Message

Received %s heartbeat from the Service Processor (SP).

### Parameters

**type** (STRING): Heartbeat from Service Processor can be of the following types: 1. IPMI Heartbeat 2. Packetized heartbeat. This parameter specifies which type of heartbeat this event corresponds to.

## sp.heartbeat.stopped

### Severity

ERROR

### Description

This event is generated when Data ONTAP does not receive Service Processor (SP) heartbeat notifications. The Service Processor and Data ONTAP exchange heartbeat messages so that they can detect when one or the other is unavailable. This event is generated when Data ONTAP has not received an expected heartbeat message from the Service Processor.

### Corrective Action

1. Connect to the SP CLI and issue: 1.1. sp version 1.2. priv set advanced 1.3. sp log debug 1.4. sp log messages 2. Run SP system diagnostics. 3. Consult Service Processor Troubleshooting Guide.

### Syslog Message

Have not received a %s heartbeat from the Service Processor (SP) in last %d seconds.

### Parameters

**type** (STRING): Heartbeat from Service Processor can be of the following types: 1. IPMI Heartbeat 2. Packetized heartbeat. This parameter specifies which type of heartbeat this event corresponds to.

**time** (INT): Heartbeats have not been received from the Service Processor (SP) for this period of time in seconds.

# sp.ipmi events

## sp.ipmi.lost.shutdown

### Severity

EMERGENCY

### Description

This message occurs when Data ONTAP® is no longer receiving the heartbeats from the Service Processor (SP) after all the SP recovery actions have been taken. Without the SP, Data ONTAP cannot monitor the health of the hardware. To prevent hardware damage and data loss, the system will shut down.

### Corrective Action

Power-cycle the system by pulling the controller out from the chassis, pushing it back in, then powering it on. If the problem persists, replace the controller module.

### Syslog Message

SP heartbeat stopped and cannot be recovered. To prevent hardware damage and data loss, the system will shut down in %d minutes.

### Parameters

**num\_minutes** (INT): Number of minutes until the system shuts down.

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

# sp.log events

## sp.log.partition.full

### Severity

ERROR

### Description

This message occurs when the available memory of the log partition in the Service Processor (SP) or the Baseboard Management Controller (BMC) is low. If the log partition is full, it might lead to unexpected configuration issues.

### Corrective Action

1. Run the "system autosupport invoke -node <node-name> -type all" command from the ONTAP CLI to collect the SP or BMC logs before they can be cleared. Verify from the command output that the logs were transferred. 2. During the next maintenance cycle, execute the following commands from the SP or BMC CLI to clear the log partition: a. "priv set diag" b. "bmc wipeclean" for BMC or "sp wipeclean" for SP c. "bmc reboot" for BMC or "sp reboot" for SP The SP or BMC "wipeclean" command clears all the configuration and logs from the SP or BMC. After you reboot, wait for 10 minutes to allow ONTAP to resync the configuration before performing any action.

### Syslog Message

%s Log partition is almost full.



## Parameters

**sp\_type** (STRING): Service processor type (SP or BMC)

## sp.log.partition.readonly

### Severity

ALERT

### Description

This message occurs when the node's Service Processor (SP) or Baseboard Management Controller (BMC) log partition is mounted as read-only. Recovery steps already have been tried automatically by the firmware, but the recovery efforts were not successful. If the log partition is mounted as read-only, the BMC or SP will not be able to record logs, configuration updates from ONTAP will fail, and ONTAP might not display the current BMC or SP status.

### Corrective Action

Search the support site knowledgebase for this event ID.

### Syslog Message

Log partition "%s" could not be recovered and is mounted as read-only.

## Parameters

**sp\_type** (STRING): Service processor type (SP or BMC)

## sp.log.partition.readwrite

### Severity

NOTICE

### Description

This message occurs when the node's Service Processor (SP) or Baseboard Management Controller (BMC) log partition is remounted as read-write after a SP or BMC reboot to recover from a log partition in a read-only state.

### Corrective Action

(None).

### Syslog Message

Log partition "%s" has been remounted as read-write after recovery from a read-only state.

## Parameters

**sp\_type** (STRING): Service processor type (SP or BMC)

## sp.network.events

### sp.network.ipv6.disable

### Severity

ERROR

**Description**

This message occurs when either a user or the system disables the network IPv6 option and the service processor (SP) is configured only with an IPv6 address.

**Corrective Action**

Ensure that the service processor (SP) is configured with an IPv4 address prior to disabling the network IPv6 option.

**Syslog Message**

The IPv6 address can be disabled only if the IPv4 address is enabled.

**Parameters**

(None).

**sp.network.ipv6.notifdisable****Severity**

NOTICE

**Description**

This message occurs when the IPv6 address of Service Processor (SP) is disabled as network IPv6 option of the system is disabled either by user or by cluster configuration.

**Corrective Action**

If IPv6 usage is desired, enable network IPv6 option of the system and assign IPv6 address to SP.

**Syslog Message**

The IPv6 address of service processor SP is disabled.

**Parameters**

(None).

**sp.network.link.down****Severity**

ERROR

**Description**

This message occurs when the Service Processor (SP) detects a link error on the SP network port. This can happen if a network cable is not plugged into the SP network port. It can also happen if the network that the SP is connected to cannot run at 10/100 Mbps.

**Corrective Action**

1. Check whether the network cable is correctly plugged into the SP network port.
2. Check the link status LED on the SP.
3. Verify that the network that the SP is connected to supports autonegotiation to 10/100 Mbps or is running at one of those speeds; otherwise, SP network connectivity does not work.

**Syslog Message**

Service Processor (SP) network port link down due to cable or network errors.

## Parameters

(None).

# sp.notconfigured events

## sp.notConfigured

### Severity

ERROR

### Description

This message occurs weekly to remind you to configure the Service Processor (SP). The SP is a physical device that is incorporated into your system to provide remote access and remote management capabilities. To use the full functionality of the SP, you need to configure it first.

### Corrective Action

1. Use the 'system service-processor network modify' command to configure the SP. If necessary, use the 'system service-processor network show' command to obtain its MAC address. 2. Use the 'system service-processor network show' command to verify the SP network configuration. 3. Use the 'system service-processor autosupport invoke' command to verify that the SP can send AutoSupport email. Note that AutoSupport mailhosts and recipients must be properly configured in Data ONTAP® before issuing this command.

### Syslog Message

The system's Service Processor (SP) is not configured. Use the 'system service-processor network modify' command to configure it.

### Parameters

**object\_type** (STRING): Type of resource object under notification. For this EMS, the object\_type will always be NODE.

**object\_uuid** (STRING): UUID of the resource object. For this EMS, the UUID will be of the node.

# sp.orftp events

## sp.orftp.failed

### Severity

NOTICE

### Description

This message occurs when there is a communication error while sending information to or receiving information from the Service Processor (SP). This error could be due to a communication error while the information is being sent or received, or the SP might not be operational.

### Corrective Action

1. Run the 'system service-processor show' command at the Data ONTAP prompt to check whether the SP is operational. 2. If the SP is operational and this message persists, run the 'system service-processor reboot-sp' command at the Data ONTAP prompt to reboot the SP. 3. If this message persists after the SP has been rebooted, contact NetApp technical support.

## Syslog Message

SP communication error, %s.

## Parameters

**reason** (STRING): Reason for the failure.

# sp.postupd events

## sp.postupd.ck.unexpt.sp.rev

### Severity

ERROR

### Description

This message occurs after a scheduled firmware update, when the servprocd daemon detects that the Service Processor(SP) firmware version changed to an unexpected value.

### Corrective Action

Check whether the SP firmware has been updated through the SP CLI. If not: 1. Manually update the SP firmware. 2. If the failures persist, contact NetApp technical support.

## Syslog Message

%s.

## Parameters

**reason** (STRING): Reason for the failure: %s.

# sp.reboot events

## sp.reboot.ro.partition

### Severity

NOTICE

### Description

This message occurs when the node's Service Processor (SP) or Baseboard Management Controller (BMC) log partition is mounted as read-only after an internal error has occurred. The system automatically reboots the SP or BMC to attempt to remount the log partition as read-write.

### Corrective Action

(None).

## Syslog Message

Rebooting the %s because the log partition is in a read-only state.

## Parameters

**sp\_type** (STRING): Service processor type (SP or BMC)

## sp.reboot.sensor.unreadable

### Severity

NOTICE

### Description

This message occurs when one or more sensors are unreadable from Service Processor (SP) or Baseboard Management Controller (BMC). SP or BMC is rebooted in attempt to recover the sensor reading.

### Corrective Action

No user action is required.

### Syslog Message

Rebooting %s because one or more sensors are unreadable.

### Parameters

**sp\_type** (STRING): Service processor type (SP or BMC).

## sp.servprocd events

### sp.servprocd.upd.error

### Severity

ERROR

### Description

This message occurs in the following cases: 1. Scheduling of Service Processor (SP) firmware update failed. 2. SP firmware update failed after it has been scheduled. 3. Post-update checks have failed.

### Corrective Action

1. Perform a manual SP firmware update. 2. If the failures persist, contact NetApp technical support.

### Syslog Message

SP update error: %s.

### Parameters

**reason** (STRING): SP update error: %s.

## sp.snmp events

### sp.snmp.traps.off

### Severity

INFORMATIONAL

### Description

This message occurs each time a system boots, if the advanced privilege level in Data ONTAP® was used to disable the SNMP Trap feature of the Service Processor (SP). This message also occurs when the SNMP Trap capability is disabled and a user invokes a Data ONTAP command to use the SP to send an SNMP trap.

**Corrective Action**

SP SNMP Trap support is currently disabled. To enable this feature, set the 'sp.snmp.traps' option to On.

**Syslog Message**

SP's SNMP Trap support is currently disabled. Set option 'sp.snmp.traps' to 'On' to enable this feature.

**Parameters**

(None).

## sp.ssh events

**sp.ssh.access.resetup****Severity**

ERROR

**Description**

This message occurs when the SSH access list of the service processor needs to be reconfigured.

**Corrective Action**

Use the "system service-processor ssh add-allowed-addresses" command to reconfigure the service processor SSH access list.

**Syslog Message**

Service Processor SSH access list needs to be reconfigured.

**Parameters**

(None).

## sp.upd events

**sp.upd.bad.fw.package****Severity**

ERROR

**Description**

This message occurs when the Service Processor (SP) daemon cannot extract the SP firmware version from the metadata file. Either the file cannot be opened, or the number cannot be found within the file.

**Corrective Action**

Use the "system node image package delete" command to remove the corrupted package. If failures persist, contact NetApp technical support.

**Syslog Message**

Service Processor daemon failed to read metadata from the image package.

**Parameters**

**Package\_Name** (STRING): Name of the SP firmware package being processed.

## sp.upd.fw.depot.full

### Severity

ERROR

### Description

This message occurs when a user downloads more than the supported number of Service Processor (SP) firmware packages to the local node.

### Corrective Action

Use the "system node image package delete" command to remove older packages. If failures persist, contact NetApp technical support.

### Syslog Message

Repository for SP image is full.

### Parameters

**available\_sp\_pkg\_count** (INT): Number of SP firmware packages existing in the local node.

**max\_num\_of\_sp\_pkgs** (INT): Maximum number of supported SP packages in the local node.

## sp.upd.unzip.event

### Severity

ERROR

### Description

This message occurs when the Service Processor (SP) daemon cannot expand the compressed SP firmware packages for the local node. The expected package could either not be found, or it could not be expanded.

### Corrective Action

Use the "system node image package delete" command to remove the corrupted package. If failures persist, contact NetApp technical support.

### Syslog Message

Service Processor failed to decompress the image package.

### Parameters

**package\_name** (STRING): SP firmware package name.

**sp\_fw\_update\_type** (STRING): Type of SP firmware update being performed when the error was encountered.

## sp.userlist events

### sp.userlist.update.failed

### Severity

ERROR

## Description

This message occurs when there is an error updating user information for the Service Processor (SP). When user information is updated on Data ONTAP®, the SP is also updated with the new changes. This enables users to log in to the SP. User information update for the SP might have failed due to a communication error with the SP or the SP might not be operational.

## Corrective Action

1. Run the 'system service-processor show' command at the Data ONTAP prompt to check whether the SP is operational.
2. If the SP is operational and this message persists, run the 'system service-processor reboot-sp' command at the Data ONTAP prompt to reboot the SP.
3. Retry the operation that caused the error message.
4. If this message persists after the SP has been rebooted, contact NetApp technical support.

## Syslog Message

Error updating SP user information, %s.

## Parameters

**reason** (STRING): Reason for the failure.

# sp.wdog events

## sp.wdog.set.fail

### Severity

ERROR

### Description

This message occurs when ONTAP® software fails to disable or enable the service-processor' (SP or BMC) watchdog timer. This might occur due to an intermittent SP or BMC firmware failure.

### Corrective Action

If the system continues to run normally, you can ignore this message. If a system disruption occurs after this failure, reboot the SP or BMC or power-cycle the system before starting ONTAP or performing a giveback. If system disruptions continue to occur, contact NetApp technical support for assistance.

### Syslog Message

Unable to %s service-processor's watchdog timer. Error code: 0x%x.

### Parameters

**op** (STRING): Whether the failure was to enable or disable the timer.

**errcode** (INT): Error code for debugging.



# space events

## space.ratchet events

### space.ratchet.failure

**Severity**

ERROR

**Description**

This event is generated if the space ratchet catches a space-check error.

**Corrective Action**

(None).

**Syslog Message**

%s %s %s %s %s %s %s %s %s %s %s %s

**Parameters**

- f2** (STRING): Third frame of backtrace.
- f3** (STRING): Fourth frame of backtrace.
- f4** (STRING): Fifth frame of backtrace.
- f5** (STRING): Sixth frame of backtrace.
- f6** (STRING): Seventh frame of backtrace.
- f7** (STRING): Eighth frame of backtrace.
- f8** (STRING): Ninth frame of backtrace.
- f9** (STRING): Tenth frame of backtrace.
- f10** (STRING): Eleventh frame of backtrace.
- f11** (STRING): Twelfth frame of backtrace.
- f12** (STRING): Thirteenth frame of backtrace.

### space.ratchet.failure.detail

**Severity**

ERROR

**Description**

This event is generated if the space ratchet catches a space-check error.

**Corrective Action**

(None).

**Syslog Message**

%s %s

**Parameters**

- description** (STRING): Description of failure.
- buffers** (STRING): Details of buffers that were predirtied.

# spider events

## spider.message events

### spider.message.fail

#### Severity

ALERT

#### Description

This message occurs when a directory traversal task cannot finish because of an unexpected error.

#### Corrective Action

Ensure all nodes in the cluster are online and healthy, and that all volumes needed for the directory traversal are online. Then retry the operation.

#### Syslog Message

Volume %llu failed to perform directory task %s because of %s.

#### Parameters

**volume\_dsid** (LONGINT): DSID of the volume that encountered the error.

**task** (STRING): Description of the task that could not be performed.

**error** (STRING): Description of the error that occurred.

# splitter events

## splitter.owm events

### splitter.owm.recycle.start

#### Severity

NOTICE

#### Description

This message occurs when overlapping write manager recycling is activated as the allocation exceeds the threshold. This event is dependent on workload and resource management. It is resolved automatically internally.

#### Corrective Action

(None).

#### Syslog Message

Overlapping write manager recycling is activated as the allocation exceeds the threshold for relationship with UUID %s.

#### Parameters

**relationship\_id** (STRING): SnapMirror relationship universal unique identifier (UUID).

# splog events

## splog.failed events

### splog.failed.to.transfer.log

#### Severity

NOTICE

#### Description

This message occurs when the system fails to transfer the SP log file from the remote Service Processor (SP). This prevents abnormal status of the system from being diagnosed by using the SP log files.

#### Corrective Action

1. Check whether you can transfer SP log files from the remote SP manually. From the clustershell prompt: a. Enter the 'system node autosupport invoke-splog -remote-node "nodename"' command. b. Check the output message of the command to determine whether the command succeeded. 2. If you cannot transfer SP log files from the remote SP manually, then restart the SP. From the clustershell prompt, run the 'set diag' and 'system node run-service-processor -node "down\_node\_name"' commands to log in to the remote SP. From the SP prompt, run the 'sp reboot' command. Example:  
splog-cmode::> set diag splog-cmode::\*> system node run-service-processor -node splog-cmode-01 SP  
splog-cmode> sp reboot 3. After the SP finishes rebooting, if you still cannot transfer SP log files from the remote SP manually, then contact NetApp technical support for further assistance.

#### Syslog Message

Failed to transfer log from Service Processor (SP) on the remote node (%s) because of %s.

#### Parameters

**node\_name** (STRING): Name of the remote node that is down and whose SP the log file cannot be transferred from.

**reason** (STRING): Reason for the warning.

## splog.running events

### splog.running.normally

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Data ONTAP® process (splogd) that monitors the Service Processor (SP) for available SP logs starts successfully or when it returns to normal operation from degraded mode.

#### Corrective Action

(None).

#### Syslog Message

Process splogd is operating normally.

## Parameters

(None).

# splog.sp events

## splog.sp.not.accessible

### Severity

ERROR

### Description

This message occurs when the local or remote node Service Processor (SP) cannot be accessed. This prevents the SP log collection job from communicating with the SP and transferring SP log files when new log files are available. SP log files are usually used for diagnosing abnormal status of the system, for example, abnormal shutdown.

### Corrective Action

1. Check whether the SP is pingable by using the 'ping sp-ip-address' command at the clustershell prompt.
2. If the remote node SP is not pingable: a) Check whether the network cable is connected with the remote SP. b) If the remote node network cable is connected, check whether the IP address of the remote SP and the e0M port are in the same subnet. From the clustershell prompt, run the 'set diag' and 'system node run-service-processor -node node\_name' commands to log in to the remote SP. From the SP prompt, run the 'sp status' command. From the clustershell prompt, run the 'network interface show' command to get e0M port details. c) If the IP address of the remote down node SP and the e0M port are in the same subnet and the remote SP is not pingable, then restart the SP. From the clustershell prompt, run the 'set diag' and 'system node run-service-processor -node node\_name' commands to log in to the remote SP. From the SP prompt, run the 'sp reboot' command.
3. If the local node SP is not pingable: a) Check whether the IP address of the local SP and the e0M port are in the same subnet. From the clustershell prompt, press Ctrl-g to log in to the local SP. From the SP prompt, run the 'sp status' command. From the clustershell prompt, run the 'network interface show' command to get e0M port details. b) If the address of local node SP and the e0M port are in the same subnet and the SP is not pingable, then restart the SP. From the clustershell prompt, run the 'system service-processor reboot-sp -image primary' command.
4. After the SP finishes rebooting, if you still cannot ping the SP, then contact NetApp technical support for further assistance.

### Syslog Message

The SP on the remote node (%s) is not accessible because of %s.

### Parameters

**node\_name** (STRING): Name of the node whose SP is not accessible.

**reason** (STRING): Reason for the warning.

## splog.sp.not.configured

### Severity

ERROR

### Description

This message occurs when the IP address of the Service Processor (SP) is not configured properly. This prevents the SP log collection job from communicating with the SP and transferring SP log files when new log files are available. SP log files are usually used for diagnosing abnormal status of the system, for example, abnormal shutdown.

## Corrective Action

1. Check the IP address of the SP. Remote node SP case: From the clustershell prompt, run the 'set diag' and 'system node run-service-processor -node node\_name' commands to log in to the remote SP. From the SP prompt, run the 'sp status' command. Local node SP case: From the clustershell prompt, press 'Ctrl-g' and log into SP. From the SP prompt, run the 'sp status' command. 2. If there is no IP address for the SP, or if the IP address is invalid, then restart the node. The IP address of the SP will be configured by the host software. 3. If there is still no IP address for the SP after step 2, then contact NetApp technical support for further assistance.

## Syslog Message

The SP on the node (%s) is not configured properly.

## Parameters

**node\_name** (STRING): Name of the node whose SP has its IP address configured incorrectly.

# splog.unsupported events

## splog.unsupported

### Severity

INFORMATIONAL

### Description

This message occurs when the current firmware of the Service Processor (SP) does not support the log collection function.

### Corrective Action

Check the firmware version from the SP CLI using the 'version' command and compare it to the most current supported version for your storage system. If necessary, download and update the SP firmware to the latest supported version.

### Syslog Message

The firmware on the SP does not support the log collection function.

### Parameters

(None).

# splog.warnings events

## splog.warnings

### Severity

ERROR

### Description

This message occurs when the ONTAP® splogd process has been unable to retrieve a log from the Service Processor (SP) or baseboard management controller (BMC) for at least 24 hours.

### Corrective Action

1. Use the "system service-processor show -node <node-name>" command to query the status of the SP

or BMC. Contact NetApp technical support for further assistance if the SP or BMC is not operational. 2. If the status of the SP or BMC is "online," then attempt to restart log monitoring by using the "system node autosupport invoke-splog -remote-node <node-name>" command. Verify from the command output that the logs were transferred successfully. (This command might require up to three attempts.) 3. If every attempt has failed, then restart the SP or BMC with the "system service-processor reboot-sp -node <node-name>" command. 4. After the SP or BMC has finished rebooting, repeat Step 2. 5. If the SP or BMC is still not transferring logs properly, run the "sp status -d" command from the SP CLI prompt or the "bmc status -d" command from the BMC CLI prompt, and include the resulting output when contacting NetApp technical support for further assistance.

### **Syslog Message**

Could not retrieve logs from the %s for 24 hours: %s

### **Parameters**

**sp\_type** (STRING): SP type

**reason** (STRING): Reason for the warning.

# spm events

## spm.avs events

### spm.avs.process.exit

#### Severity

ERROR

#### Description

This message occurs when the antivirus service stops running due to repeated failures. This can be the result of a bad or corrupted update in the scan engine library or virus definition files, or due to an internal failure. The failure of this subsystem can affect client access if mandatory scanning is enabled. It could also result in files not being scanned.

#### Corrective Action

Disable and then reenble the antivirus engine. If the issue persists, contact NetApp technical support.

#### Syslog Message

Antivirus server with ID %d exited as a result of signal %s. The service will attempt to restart.

#### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.bgpd events

### spm.bgpd.process.exit

#### Severity

NOTICE

#### Description

This message occurs when the Border Gateway Protocol daemon (bgpd) involuntarily terminates. This daemon is responsible for communicating route information between a storage system and gateway for the virtual IP (VIP) feature. This failure will be automatically addressed by spmctl restarting bgpd. However, if bgpd fails to restart for any reason, the VIP feature will not function correctly.

#### Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport@ message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

#### Syslog Message

Border Gateway Protocol daemon (bgpd) with ID %d exited as a result of signal: %s. The service will attempt to restart.

#### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Cause of the process exit.



# spm.coresegd events

## spm.coresegd.process.exit

### Severity

ALERT

### Description

This message occurs when the core segment server (coresegd) involuntarily terminates. The termination of this server impacts all core segmenting services provided by the node.

### Corrective Action

The system attempts to recover by restarting the node connectivity "process server". If the process does not restart within the threshold number of retries, then an AutoSupport(tm) message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

The core segment server (coresegd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process terminated.

# spm.cphmd events

## spm.cphmd.process.exit

### Severity

ALERT

### Description

This message occurs when the Chassis Health Monitor process (cphmd) involuntarily terminates. The termination of this process can affect monitoring of the subsystems of CPHM .

### Corrective Action

The system attempts to recover by restarting the chassis HM process. server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

Controller Health Monitor process (cphmd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.crs events

## spm.crs.process.exit

### Severity

ALERT

### Description

This message occurs when the configuration replication server (crs) involuntarily terminates.

### Corrective Action

The system attempts to recover by restarting the configuration replication process server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

The configuration replication server (crs) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.cshmd events

### spm.cshmd.process.exit

### Severity

ALERT

### Description

This message occurs when the Cluster Switch Health Monitor process (cshmd) involuntarily terminates. The termination of this process can affect monitoring of the subsystems of CSHM .

### Corrective Action

The system attempts to recover by restarting the cluster switch HM process. server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

Cluster Switch Health Monitor process (cshmd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.enclmgmtd events

### spm.enclmgmtd.process.exit

## Severity

NOTICE

## Description

This message occurs when the Enclosure Management Daemon (enclmgmtd) involuntarily terminates. This daemon is responsible for enclosure monitoring for this platform. This failure does not affect client access. However, enclosure monitoring might be in a degraded state if the service fails to restart.

## Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport® message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

## Syslog Message

Enclosure Management Daemon (enclmgmtd) with ID %d exited as a result of signal %s. The service will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.envmgr events

## spm.envmgr.process.exit

### Severity

ERROR

### Description

This message occurs when the environmental monitoring daemon (env\_mgr) involuntarily terminates. This daemon is responsible for environmental monitoring for this platform. This failure does not affect client access; however, environmental monitoring might be in a degraded state if the service fails to restart, which could result in an unpredicted shutdown if an adverse environmental condition or hardware failure occurs.

### Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, call technical support.

### Syslog Message

Environmental monitoring daemon (env\_mgr) with ID %d exited as a result of signal %s. The service will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.fpolicy events

## spm.fpolicy.process.exit

### Severity

ALERT

### Description

This message occurs when the fpolicy involuntarily terminates. This can be due to one or more causes and might be related to some recent or previous operations. This subsystem is responsible for sending notification to external fpolicy servers and failure of this subsystem can affect client access.

### Corrective Action

While this problem persists, consider migrating the Logical Interfaces (LIFs) hosted on the current node to another node in the cluster and ensure that the auto-revert and DNS load-balancing features do not migrate the LIFs back to the current node. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Fpolicy Service Manager process with ID %d exited as a result of signal %s. The service will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.hashd events

### spm.hashd.process.exit

#### Severity

ALERT

#### Description

This message occurs when the hashd server (hashd) involuntarily terminates.

#### Corrective Action

The system attempts to recover by restarting the hashd process server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

#### Syslog Message

The hashd server (hashd) with ID %d exited as a result of %s. This server will attempt to restart.

#### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.httpd events

## spm.httpd.process.exit

### Severity

ALERT

### Description

This message occurs when the web server (httpd) involuntarily terminates. The termination of this server will impact all HTTP/HTTPS services provided by the node.

### Corrective Action

The system attempts to recover by restarting the node connectivity process. server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

The web server (httpd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.ktlsd events

## spm.ktlsd.process.exit

### Severity

ALERT

### Description

This message occurs when the Kernel Transport Layer Security Daemon support process (ktlsd) involuntarily terminates. The termination of this process can affect monitoring of the subsystems of KTLS.

### Corrective Action

The system attempts to recover by restarting the KTLS Daemon process. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

Kernel TLS support process (ktlsd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.mes events

## spm.mes.process.exit

### Severity

ERROR

### Description

This message occurs when the SPD MES daemon (spd) involuntarily terminates. This daemon is responsible for updating DIMM SPDs with correctable error statistics. This failure doesn't affect client access.

### Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, call technical support.

### Syslog Message

SPD MES daemon (spd) with ID %d exited as a result of signal %s. The service will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.mgwd events

### spm.mgwd.process.exit

### Severity

EMERGENCY

### Description

This message occurs when the Management Gateway (mgwd) subsystem involuntarily exits. The subsystem provides administration services to manage the cluster and current node. This failure can disrupt administrative tasks that are being performed on the current node. The system attempts to recover by restarting the subsystem. While this situation persists, if a Cluster Management Logical Interface (LIF) is hosted on the current node, migrate it to another node in the cluster using another node management LIF or console access to another node. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent.

### Corrective Action

In the rare event that the Management Gateway (mgwd) subsystem has terminated unexpectedly, the automatic restart of the process will often stabilize the situation. While this situation persists, if a Cluster Management Logical Interface (LIF) is hosted on the current node, migrate it to another node in the cluster using another node management LIF or console access to another node. In some cases, there is a systemic issue that can be cleared up by rebooting the node when multiple restarts do not stabilize mgwd. This might be done automatically by the Node Watchdog service for persistent issues within a short window of time. In other cases, recent changes in logging, configuration or external management activities can be contributing factors to the unexpected terminations. Consider reverting one or more of these changes. If the issue persists, contact NetApp technical support.

### Syslog Message

Management Gateway (mgwd) subsystem with ID %d exited as a result of signal %s. The subsystem will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.mntsvc events

## spm.mntsvc.process.exit

### Severity

ALERT

### Description

This message occurs when the Node Root Mount Service (mntsvc) involuntarily terminates. The termination of this service affects all Remote Support Agent services provided by the node.

### Corrective Action

The system attempts to recover by restarting the process. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

The Node Root Mount Service (mntsvc) with ID %d exited as a result of %s. This server will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.named events

## spm.named.process.exit

### Severity

ERROR

### Description

This message occurs when the Domain Name Service (DNS) name resolution daemon 'named' involuntarily exits. This error can have one or more causes and might be related to some recent or previous operations. The failure of this daemon can affect DNS zone-based access. In particular, the cluster might no longer be able to resolve DNS zone names to IP addresses and clients that connect to the zone using zone names.

### Corrective Action

While this problem persists, ensure that any new mounts use an IP address instead of a zone name. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

DNS name resolution daemon 'named' with ID %d exited as a result of signal %s. The service will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.nchmd events

## spm.nchmd.process.exit

### Severity

ALERT

### Description

This message occurs when the Node Connectivity Health Monitor process (nchmd) involuntarily terminates. The termination of this process can affect node connectivity monitoring for SAS subsystem.

### Corrective Action

The system attempts to recover by restarting the node connectivity process. server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

Node Connectivity Health Monitor process (nchmd) with ID %d exited as a result of %s. This server will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.ndmpd events

## spm.ndmpd.process.exit

### Severity

ALERT

### Description

This message occurs when the Network Data Management Protocol (NDMP) server (ndmpd) involuntarily terminates. The termination of this server can affect NDMP client access.

### Corrective Action

The system attempts to recover by restarting the NDMP server. If the NDMP server does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Network Data Management Protocol server (ndmpd) with ID %d exited as a result of %s. This server will attempt to restart.

## Parameters

**id** (INT): Process identifier.



**cause** (STRING): Plain text description of why the process exited.

## spm.nphmd events

### spm.nphmd.process.exit

#### Severity

ALERT

#### Description

This message occurs when the Controller Health Monitor process (nphmd) involuntarily terminates. The termination of this process can affect monitoring of the subsystems of NPHM .

#### Corrective Action

The system attempts to recover by restarting the controller HM process. server. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

#### Syslog Message

Controller Health Monitor process (nphmd) with ID %d exited as a result of %s. This server will attempt to restart.

#### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

## spm.ntpd events

### spm.ntpd.process.exit

#### Severity

ERROR

#### Description

This message occurs when the Network Time Protocol Daemon process (ntpd) involuntarily terminates. The termination of this process can affect synchronization of system time with network time servers.

#### Corrective Action

The system attempts to recover by restarting the ntpd process. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

#### Syslog Message

The ntpd with ID %d exited as a result of %s. This server will attempt to restart.

#### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Reason the process exited.

# spm.perfstatd events

## spm.perfstatd.process.exit

### Severity

ALERT

### Description

This message occurs when the perfstat daemon (perfstatd) involuntarily terminates.

### Corrective Action

The system attempts to recover by restarting the perfstat daemon. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

The perfstat daemon(perfstatd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.pipd events

## spm.pipd.process.exit

### Deprecated

Deprecated as of Data ONTAP 9.7 because the pipd process no longer exists.

### Severity

ALERT

### Description

This message occurs when the post-ingest processing daemon (pipd) involuntarily terminates. The termination of this service prevents data policies and constituent balancing from being applied to Infinite Volume data owned by the node in question.

### Corrective Action

The system attempts to recover by restarting the pipd. If the process does not restart within the threshold number of retries, then an AutoSupport(tm) message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

The pipd with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.schmd events

## spm.schmd.process.exit

### Severity

ALERT

### Description

This message occurs when the System Connectivity Health Monitor process (schmd) involuntarily terminates. The termination of this process can affect system connectivity monitoring for HA pair in SAS subsystem.

### Corrective Action

The system attempts to recover by restarting the system connectivity process. If the process does not restart within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact technical support.

### Syslog Message

System Connectivity Health Monitor process (schmd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.secd events

## spm.secd.process.exit

### Severity

EMERGENCY

### Description

This message occurs when the Security daemon (secd) involuntarily terminates. This can be due to one or more causes and might be related to some recent or previous operations. This subsystem is responsible for client authentication and failure of this subsystem can affect client access.

### Corrective Action

While this problem persists, consider migrating the Logical Interfaces (LIFs) hosted on the current node to another node in the cluster and ensure that the auto-revert and DNS load-balancing features do not migrate the LIFs back to the current node. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Security daemon with ID %d exited as a result of signal %s. The service will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.servprocd events

## spm.servprocd.process.exit

### Severity

ALERT

### Description

This message occurs when the Service processor daemon process (servprocd) involuntarily terminates.

### Corrective Action

The system attempts to recover by restarting the web server. If the server does not restart within the threshold number of retries, then an AutoSupport(tm) message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Service processor daemon (servprocd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.shmd events

## spm.shmd.process.exit

### Severity

ALERT

### Description

This message occurs when the System Health Monitor process (shmd) involuntarily terminates. The termination of this process can affect the system health management and diagnosis feature.

### Corrective Action

The system attempts to recover by restarting the shmd process. If the server does not restart within the threshold number of retries, then an AutoSupport(tm) message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

System Health Monitor process (shmd) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.timestate events

## spm.timeState.process.exit

### Deprecated

This event is removed since time\_state server was removed in 8.3.

### Severity

ERROR

### Description

This message occurs when the subsystem managing timekeeping parameters (time\_state) stops running. The failure of this subsystem can cause Network Time Protocol (NTP) settings and the time zone to be inconsistent among subsystems. This might also have an effect on time displayed in log messages.

### Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Subsystem managing timekeeping parameters (time\_state) with ID %d exited as a result of signal %s. The subsystem will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process stopped running.

## spm.ucoreman events

### spm.ucoreman.process.exit

### Severity

NOTICE

### Description

This message occurs when the Usercore Manager (ucoreman) daemon involuntarily terminates. This daemon is responsible for generating a crash report when an application crashes and dumps core. The crash report helps technical support to diagnose issues with the application. This failure does not affect client access.

### Corrective Action

The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, call technical support.

### Syslog Message

Usercore Manager (ucoreman) with ID %d exited as a result of signal %s. The service will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.upgrademgr events

## spm.upgrademgr.process.exit

### Severity

ERROR

### Description

This message occurs when the upgrade manager server (upgrademgr) involuntarily terminates.

### Corrective Action

The system attempts to recover by restarting the upgrademgr process server. If the process does not restart within the threshold number of retries, then an AutoSupport® message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

The upgrade manager server (upgrademgr) with ID %d exited as a result of %s. This server will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.vifmgr events

## spm.vifmgr.process.exit

### Severity

EMERGENCY

### Description

This message occurs when the Logical Interface Manager (VIFMgr) subsystem aborts. This error can have one or more causes and might be related to recent or previous operations. The failure of this subsystem can result in logical interfaces (LIFs) on the current node migrating to another node within the same cluster.

### Corrective Action

Ensure that the LIFs on the current node are hosted on another node in the cluster; if they are not, migrate them to another node. While this problem persists, you should also ensure that the auto-revert and DNS load-balancing features do not migrate the LIFs back to the current node. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Logical Interface Manager(VifMgr) with ID %d aborted as a result of signal %s. The subsystem will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process aborted.

# spm.vldb events

## spm.vldb.process.exit

### Severity

EMERGENCY

### Description

This message occurs when Volume Location Database (VLDB) subsystem stops running. This can be due to one or more causes and might be related to some recent or previous operations. The failure of this subsystem can affect client access.

### Corrective Action

Ensure that none of the volume/aggregate operations are performed on the current node. While this problem persists, you might want to migrate the Logical Interfaces (LIFs) to another node in the cluster and ensure that auto-revert and DNS load-balancing features do not migrate the LIFs back to the current node. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

Volume Location Database(VLDB) subsystem with ID %d exited as a result of signal %s. The subsystem will attempt to restart.

### Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.vserverdr events

## spm.vserverdr.process.exit

### Severity

ALERT

### Description

This message occurs when the Vserver DR process involuntarily terminates. This might cause interruptions in the automatic or scheduled configuration backup of the Vserver's data and configuration.

### Corrective Action

The system attempts to recover by restarting the Vserver DR process. Use the "debug process show -process-name vserverdr" command in the diagnostic privilege level to check whether the Vserver DR process has successfully restarted. If the process does not restart within the threshold number of retries, an AutoSupport(tm) message will be sent. If your system is not configured to send AutoSupport messages and the process has not automatically recovered, contact NetApp technical support.

### Syslog Message

The Vserver DR process with ID %d exited as a result of %s. The system will attempt to restart it.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Plain text description of why the process exited.

# spm.yplibind events

## spm.yplibind.process.exit

### Severity

ALERT

### Description

This message occurs when the NIS client binding daemon 'ypbind' involuntarily exits. This error can have one or more causes and might be related to some recent or previous operations. The termination of this daemon will affect all NIS-based services provided by the controller. In particular, it will no longer be possible to use NIS as a name service source for name lookups.

### Corrective Action

While this problem persists, unconfigure NIS as a name service source in the name service switch database. For passwd, group, or netgroup lookups, use an alternative name service source, such as LDAP or files. The system attempts to recover by restarting the subsystem. If the subsystem does not recover within the threshold number of retries, then an AutoSupport(tm) message is sent. If your system is not configured to send AutoSupport messages and the issue persists, contact NetApp technical support.

### Syslog Message

NIS client binding daemon 'ypbind' with ID %d exited as a result of signal %s. The service will attempt to restart.

## Parameters

**id** (INT): Process identifier.

**cause** (STRING): Reason the process exited.



# spmgmt events

## spmgmt.driver events

### spmgmt.driver.hourly.stats

#### Severity

ERROR

#### Description

This message occurs when the system encounters an error while trying to get hourly statistics from the Service Processor (SP). The error could be due a communication error with the SP or the SP might not be operational.

#### Corrective Action

1. Run the 'system service-processor show' command at the Data ONTAP® prompt to check whether the SP is online. 2. If the SP is online and this message persists, run the 'system service-processor reboot-sp' command at the Data ONTAP prompt to reboot the SP.

#### Syslog Message

The software driver for the Service Processor (SP) detected a problem: %s.

#### Parameters

**reason** (STRING): Reason for this failure.

### spmgmt.driver.mailhost

#### Severity

ERROR

#### Description

This message occurs when the Service Processor (SP) setup cannot connect to the mailhost specified in Data ONTAP®.

#### Corrective Action

1. Check the system AutoSupport configuration to verify that a valid mailhost is configured in Data ONTAP.
2. Ensure that Data ONTAP can successfully connect to the specified mailhost by running a test command to invoke AutoSupport.

#### Syslog Message

SP setup could not access the mailhost specified in Data ONTAP.

#### Parameters

(None).

### spmgmt.driver.network.failure

#### Severity

ERROR

## Description

This message occurs when the system encounters a failure during network configuration of the Service Processor (SP). The system cannot assign the SP a DHCP (Dynamic Host Configuration Protocol) or fixed IP address.

## Corrective Action

1. Check whether the network cable is correctly plugged into the SP network port.
2. Check the link status LED on the SP.
3. Verify that the network that the SP is connected to has autonegotiation enabled.

## Syslog Message

The network configuration of the Service Processor (SP) failed due to cable or network errors.

## Parameters

(None).

## **spmgmt.driver.timeout**

## Severity

ERROR

## Description

This message occurs when there is a failure during communication with the Service Processor (SP) firmware. The failure could be due to a communication error with the SP or the SP might not be operational.

## Corrective Action

1. Issue the 'system service-processor show' command at the Data ONTAP® prompt to check whether the SP is online.
2. If the SP is operational and this message persists, run the 'system service-processor reboot-sp' command at the Data ONTAP prompt to reboot the SP. After the reboot, this message should no longer occur. If the message occurs again, contact NetApp technical support and explain that you already performed the preceding steps.

## Syslog Message

The software driver for the Service Processor (SP) detected a problem: %s.

## Parameters

**reason** (STRING): Reason for the failure.

# sshd events

## sshd.auth events

### sshd.auth.loginDenied

**Severity**

NOTICE

**Description**

This event is issued when sshd refuses a login attempt due to authentication failure.

**Corrective Action**

Use a valid username/password combination to login.

**Syslog Message**

(None).

**Parameters**

**message** (STRING): Details about the failed login attempt.

## sshd.logingracetime events

### sshd.loginGraceTime.expired

**Severity**

ERROR

**Description**

This message occurs when a user tries to establish a Secure Shell (SSH) connection to a storage system and does not provide the password within the allotted timeout period. Many such connection attempts could potentially disallow other users from logging in to the storage system, causing a Denial of Service (DOS) attack.

**Corrective Action**

If the remote host is retrying the SSH connection repeatedly, block the remote host by adding its IP address to the deny list using the "firewall policy" command.

**Syslog Message**

Timeout before password authentication for remote host %s.

**Parameters**

**ipaddress** (STRING): IP address of the remote host that was trying to establish an SSH connection.

# stackmon events

## stackmon.stack events

### stackmon.stack.id.conflict

#### Severity

ERROR

#### Description

This message occurs when the system cannot determine the stack ID for a disk shelf. The system has exceeded the maximum number of times it can change the stack ID of the shelf. The cause of this issue is a misconfiguration in the system.

#### Corrective Action

Run the 'system health alert show' command and resolve all of the issues reported by following the corrective actions provided.

#### Syslog Message

Stack ID of the Shelf %d (S/N %s) on channel %s cannot be agreed on by all the nodes in the cluster.

#### Parameters

**shelfIdent** (INT): ID number displayed on the disk shelf OPS panel.

**serialNumber** (STRING): Serial number of disk shelf.

**channelName** (STRING): Channel connected to the disk shelf.

# storage events

## storage.cmd events

### storage.cmd.balanceRequest

**Severity**

INFORMATIONAL

**Description**

Indicates a user request to balance I/O across initiator ports has been issued.

**Corrective Action**

(None).

**Syslog Message**

Storage related I/O has been balanced across available initiator ports.

**Parameters**

(None).

### storage.cmd.logPortIO

**Severity**

INFORMATIONAL

**Description**

Logs the mix of I/O bound to a specific initiator port prior to balancing. Issued for all initiator ports when a user requests a balancing of I/O.

**Corrective Action**

(None).

**Syslog Message**

Initiator Port %s had %d percent of all storage bound I/O traffic routed through it.

**Parameters**

**portName** (STRING): The name of the initiator port

**ioPct** (INT): The percentage of overall backend I/O routed through this port

## storage.interface events

### storage.interface.linkdown

**Severity**

ALERT

**Description**

This message occurs when a storage interface reports a link down.

**Corrective Action**

Check the system and attempt to bring the network interface back online.

**Syslog Message**

Storage interface link status is down. Node ID: %u, Node UUID: %s, Storage Interface: %s

**Parameters**

**nodeID** (INT): Unique identifier of the node that is hosting the storage interface.

**nodeUuid** (STRING): Universal unique identifier of the node that is hosting the data IP.

**sipi** (STRING): Storage Interface name

**storage.interface.linkup****Severity**

NOTICE

**Description**

This message occurs when a storage interface reports a link up.

**Corrective Action**

(None).

**Syslog Message**

Storage interface link status is up. Node ID: %u, Node UUID: %s, Storage Interface: %s

**Parameters**

**nodeID** (INT): Unique identifier of the node that is hosting the storage interface.

**nodeUuid** (STRING): Universal unique identifier of the node that is hosting the data IP.

**sipi** (STRING): Storage Interface name

# storencr events

## storencr.disk events

### storencr.disk.opDone

#### Severity

INFORMATIONAL

#### Description

This message occurs when a Storage Encryption disk command operation that modifies encrypting drive parameters is completed successfully.

#### Corrective Action

(None).

#### Syslog Message

Storage Encryption disk %s operation was completed successfully.

#### Parameters

**operation** (STRING): Name of the operation.

# storvsc events

## storvsc.srb events

### storvsc.srb.status

#### Severity

ALERT

#### Description

This message occurs when ONTAP® software receives an unexpected hypervisor I/O adapter driver response code to an I/O request. The I/O request is retried.

#### Corrective Action

If the condition persists, contact the hypervisor provider for the corrective action.

#### Syslog Message

%s. The I/O request will be retried.

#### Parameters

**Descriptor** (STRING): Hypervisor's I/O adapter driver received response code.



# suspend events

## suspend.zapi events

### suspend.zapi

#### Severity

ALERT

#### Description

This message occurs when ONTAPI (ZAPI) has been suspended due to a lack of use.

#### Corrective Action

Ensure that all ONTAP data storage management functions using ONTAPI (ZAPI) have transitioned to ONTAP REST API. To unsuspend ONTAPI (ZAPI), run the "system services web ontapi modify -suspended false" command.

#### Syslog Message

ONTAPI (ZAPI) was suspended on "%s" due to a lack of use.

#### Parameters

**date** (STRING): Date ONTAPI (ZAPI) was suspended.

# sv events

## sv.vfiler events

### SV.vfiler.config.move.error

#### Severity

ERROR

#### Description

This message occurs when movement of SnapVault® configuration from 'from' vFiler to 'to' vFiler fails. Subsequent SnapVault transfer on the vFiler might fail.

#### Corrective Action

Move the volume back to 'from' vFiler and then again perform a vFiler move from 'from' vFiler to 'to' vFiler. Please use 'vfiler move' command to move the volume between the vfilers.

#### Syslog Message

Movement of SnapVault configuration failed while moving volumes between vFiler %s to vFiler %s. Move the volume back to %s vFiler and perform the vFiler move operation again.

#### Parameters

**from\_vfiler1** (STRING): vFiler(tm) from which resources are being moved.

**to\_vfiler** (STRING): vFiler to which resources are being moved.

**from\_vfiler2** (STRING): vFiler to which resources are being moved.

## sv.vol events

### SV.vol.snapshot.rename.error

#### Severity

ERROR

#### Description

This message occurs when the renaming operation of SnapVault® Snapshot(tm) copies belonging to a certain volume fails, even though the "vFiler rename" or "vFiler move" operation was successful. The problem might be that the volume was offline during the "vFiler rename" or "vFiler move" operation, and it might cause SnapVault to have an extra Snapshot copy on the destination. To avoid this scenario, volumes need to be online during a "vFiler rename" or "vFiler move".

#### Corrective Action

If necessary, delete the extra Snapshot copy manually after the next SnapVault update.

#### Syslog Message

Rename of SnapVault snapshots for volume %s failed. This might cause SnapVault to have an extra snapshot on destination. If necessary, delete the extra Snapshot copy manually after the next SnapVault update.

## Parameters

**vol\_name** (STRING): Volume whose Snapshot copies created by SnapVault® could not be renamed.

# sync events

## sync.snapmirror events

### sync.snapmirror.arl.veto

#### Severity

ALERT

#### Description

This message occurs when a storage failover (SFO) operation is vetoed because the aggregate is hosting volumes that are in a SnapMirror® Synchronous relationship with StrictSync policy.

#### Corrective Action

Either convert all SnapMirror Synchronous relationships with StrictSync Policy to Sync Policy for the volumes in the aggregate and then retry the operation, or retry the operation with the "override-vetoes true" parameter. Note: Using this option can cause an I/O disruption.

#### Syslog Message

SnapMirror vetoed storage failover operation because aggregate "%s" is hosting volumes that are in a SnapMirror Synchronous relationship with StrictSync policy.

#### Parameters

**aggrname** (STRING): Name of the aggregate.

### sync.snapmirror.oos.frequent

#### Severity

ERROR

#### Description

This message occurs when a SnapMirror® Sync relationship frequently changes status from "in-sync" to "out-of-sync". This behaviour impacts data protection. I/O restrictions are imposed on the source volume based on the mode of replication. For relationships with policy type "strict-sync-mirror", both client read AND write access are disallowed.

#### Corrective Action

Check network connection between the source and destination volumes. Monitor the SnapMirror® Sync relationship status by using the "snapmirror show" command and observe field "last-transfer-error" to identify the reasons for frequent status change of the SnapMirror® Sync relationship to "out-of-sync".

#### Syslog Message

Source volume "%s" and destination volume "%s" with relationship UUID "%s" is frequently undergoing change in status from "in-sync" to "out-of-sync".

#### Parameters

**srcpath** (STRING): Path name of the source volume of the SnapMirror relationship.

**dstpath** (STRING): Path name of the destination volume of the SnapMirror relationship.

**relationship\_id** (STRING): Relationship UUID.

# sysconfig events

## sysconfig.sysconfigtab events

### sysconfig.sysconfigtab.openFailed

**Severity**

NOTICE

**Description**

This event is issued when the sysconfigtab file cannot be opened.

**Corrective Action**

(None).

**Syslog Message**

sysconfig: %s

**Parameters**

**reason** (STRING): Reason why it can't be opened.

# sysmgr events

## sysmgr.ecv events

### sysmgr.ecv.fetch

#### Severity

ERROR

#### Description

This message occurs when acquiring the desired Effective Cluster Version (ECV) of System Manager was not accomplished within the time limit. This node will provide the last version of System Manager instead of the desired ECV. Any operations being blocked are unblocked, such as storage giveback and cluster management logical interface (LIF) migration.

#### Corrective Action

Use the diagnostic level "system manager version show" command to see the status of the node. If this behavior persists, contact NetApp technical support.

#### Syslog Message

ECV version %s of System Manager was not acquired within %d seconds.

#### Parameters

**ecv** (STRING): Current ECV determined by the capability management (CAPM) system.  
**timeout** (INT): Timeout in seconds.

## sysmgr.gb events

### sysmgr.gb.ecv

#### Severity

NOTICE

#### Description

This message occurs when acquiring the desired Effective Cluster Version (ECV) of System Manager was not accomplished within the time limit during the first storage giveback phase. This node will provide the last version of System Manager instead of the desired ECV. This node might acquire the desired version later.

#### Corrective Action

Use the diagnostic level "system manager version show" command to see the status of the node. If this behavior persists, contact NetApp technical support.

#### Syslog Message

ECV version of System Manager was not acquired within %d seconds during the first phase of storage giveback.

#### Parameters

**timeout** (INT): Timeout in seconds.

# sysmgr.upload.events

## sysmgr.upload.bad.method

### Severity

ERROR

### Description

This message occurs when an attempt to upload a file through the System Manager web service fails because the HTTP method is not POST.

### Corrective Action

Retry the upload using the System Manager web interface.

### Syslog Message

Failed to upload a file because the HTTP method "%s" is not allowed. The method must be POST.

### Parameters

**method** (STRING): HTTP method.

## sysmgr.upload.bad.type

### Severity

ERROR

### Description

This message occurs when a file upload through the System Manager web service fails because the file name has an unsupported extension.

### Corrective Action

Upload a file with one of the allowed file name extensions using the System Manager web interface.

### Syslog Message

Failed to upload file "%s" because the extension must be one of: %s.

### Parameters

**file** (STRING): File name.

**allowed\_types** (STRING): Allowed extensions for the file name.

## sysmgr.upload.exists

### Severity

NOTICE

### Description

This message occurs when a file upload through the System Manager web service fails because the file already exists. System Manager automatically deletes the existing file and retries the upload.

### Corrective Action

(None).

### **Syslog Message**

Failed to upload file "%s" because it already exists.

### **Parameters**

**file** (STRING): File name.

## **sysmgr.upload.internal**

### **Severity**

ERROR

### **Description**

This message occurs when a file upload through the System Manager web service fails because of an internal error.

### **Corrective Action**

Retry the upload using the System Manager web interface.

### **Syslog Message**

Failed to upload file "%s" because of an internal error.

### **Parameters**

**file** (STRING): File name.

## **sysmgr.upload.no.file**

### **Severity**

ERROR

### **Description**

This message occurs when an attempt to upload a file through the System Manager web service fails because the file is not received. This might happen if the file is larger than the limit or space is consumed by other uploads.

### **Corrective Action**

If the file is larger than the maximum allowed size, then it cannot be uploaded. If other uploads are in progress, allow time for those uploads to finish, and then try again. Otherwise, contact NetApp technical support.

### **Syslog Message**

Failed to upload file "%s" due to exceeding size limit %s or insufficient space.

### **Parameters**

**file** (STRING): File name.

**max\_filesize** (STRING): Maximum allowed size for the uploaded file.

## **sysmgr.upload.no.form**

### **Severity**

ERROR



## Description

This message occurs when an attempt to upload a file through the System Manager web service fails because there is a problem with the request format. This might happen if the Content-Type HTTP header is not "multipart/form-data" or there is not a form field named "file" containing the file to upload.

## Corrective Action

Retry the upload using the System Manager web interface.

## Syslog Message

Failed to upload a file because a form field named "file" was not found in the HTTP request, or the Content-Type HTTP header is not "multipart/form-data".

## Parameters

(None).

## sysmgr.upload.success

### Severity

NOTICE

### Description

This message occurs when a file is successfully uploaded through the System Manager web service.

### Corrective Action

(None).

### Syslog Message

File "%s" has been uploaded.

### Parameters

**file** (STRING): File name.

## sysmgr.upload.unauthorized

### Severity

ERROR

### Description

This message occurs when an attempt to upload a file through the System Manager web service fails because the user is unauthorized. This could be due to a failure to authenticate the user, e.g., the password or user name was wrong, or if the user's role is not granted access to the System Manager web service.

### Corrective Action

Verify that the user name and password used to log in to System Manager are correct. Verify that the user exists in the "security login show" table with application "http" and authentication method "password". Verify that the user's role, as shown in "security login show", is granted access to the "sysmgr" web service in "vserver services web access show sysmgr".

### Syslog Message

Failed to upload file "%s" because the user is unauthorized.

## Parameters

**file** (STRING): File name.

# tape events

## tape.cmd events

### tape.cmd.chkCondErr

#### Severity

ERROR

#### Description

This message occurs when a Check Condition status from the tape drive reports an informational condition or status to the controller. The condition that generated this event is an error that occurred during execution of the command and was not cleared by retrying the request, if retries were attempted.

#### Corrective Action

This Check Condition case is an error that might be caused by an intermittent error in device hardware or a flaw in the tape media. Media errors might also suggest that the tape drive head should be cleaned following the manufacturer's directions. In cases of repeated events, evaluate the device for proper operation and possible repair or replacement. Consult vendor documentation for additional information.

#### Syslog Message

%s device %s: Check Condition: SCSI opcode %s (CDB %s): %s - %s (0x%x - 0x%x 0x%x 0x%x).

#### Parameters

**deviceType** (STRING): Type of device on which the condition occurred.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**opcode** (STRING): Definition of the SCSI opcode (command) that encountered the check condition.

**cdb** (STRING): Opcode byte of the command block that encountered the check condition.

**sSenseKey** (STRING): String containing a description of the sense key returned by the device for the failed command.

**sSenseCode** (STRING): String containing a description of the additional sense code and sense code qualifier key returned by the device for the failed command.

**iSenseKey** (INTHEX): Integer value of the sense key for the failed command.

**iASC** (INTHEX): Integer value of the sense code for the failed command.

**iASCQ** (INTHEX): Integer value of the sense code qualifier for the failed command.

**iFRU** (INTHEX): Integer value of the field-replaceable unit(FRU) for the failed command.

### tape.cmd.chkCondNullSenseDbg

#### Severity

ERROR

#### Description

This message occurs when the command to the tape device fails with a status of Check Condition, but the adapter received no sense bytes containing details about the failure. Sense bytes typically accompany command failure reports, unless bus or loop stability problems, or target device or adapter errors have interfered with the sense byte transfer. This is a diagnostic message. See the "mt diag" command (from the nodeshell in clustered environments) for the control of tape diagnostic messages.

### Corrective Action

Ensure that the cables are securely attached, and perform other maintenance typically required for SCSI or FC stability. If the error persists, contact NetApp technical support for assistance.

### Syslog Message

%s device %s reports Check Condition status on command opcode 0x%x. No sense bytes accompany the status.

### Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**opcode** (INHEX): Opcode of the failed SCSI command.

## tape.cmd.cmdFail

### Severity

ERROR

### Description

This message occurs when a tape command fails with an error detected by the adapter. This is not an error reported by the tape drive. The adapter status values reported under this message indicate that the adapter suffered an internal error, or that the bus or loop is unstable, or that the adapter has no path to the requested device.

### Corrective Action

Verify that the loop or bus cabling is correctly installed. If the error occurs only on this device and not on others on the same interface, ensure that the external and internal cabling, if any, of this device are securely connected. If the error occurs on more than one device, then a more general debug of the cabling infrastructure might be required.

### Syslog Message

Tape device %s: command failed, ha\_status 0x%x.

### Parameters

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**adapterError** (INHEX): Internal adapter error that was detected.

## tape.cmd.scsiErrNullSense

### Severity

ERROR

### Description

This message occurs when the device reports that the command failed, but the adapter received no sense bytes containing detail about the failure. Sense bytes typically accompany command failure reports, unless bus or loop stability problems, or target device or adapter errors have interfered with the sense byte transfer.

### Corrective Action

Ensure that the cables are securely attached, and perform other maintenance typically required for SCSI or FC stability. If the error persists, contact NetApp technical support for assistance.

## Syslog Message

%s device %s reports %s status on command opcode 0x%x. No sense bytes accompany the status.

## Parameters

**deviceType** (STRING): Device type.

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**scsiStatus** (STRING): Status returned by the device.

**opcode** (INHEX): Opcode of the failed SCSI command.

# tape.config events

## tape.config.badKey

### Severity

ERROR

### Description

This message occurs when the tape configuration file has a format error in a required parameter or is missing a required parameter. The error prevents the file from being used to specify parameters for the tape drive. Data ONTAP® examines tape configuration files in the storage controller's /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image. This error can occur if a configuration file is altered, an action that is not recommended.

### Corrective Action

Recover the original file from a backup file or Snapshot(tm) copy, if possible. A copy of the file might be available from NetApp Support, although not all tape configuration files are made available there.

## Syslog Message

Required key (%s) is missing or badly formatted in tape configuration file %s.

## Parameters

**badKey** (STRING): Required parameter.

**configFile** (STRING): Name of the file with the error.

## tape.config.badOpt

### Severity

INFORMATIONAL

### Description

This message occurs when Data ONTAP® detects an incorrect format on an option key in a tape configuration file. Data ONTAP will use a default setting for this key. This message can occur if a configuration file is altered, an action that is not recommended. This condition does not prevent the use of the specified tape device. Data ONTAP examines tape configuration files in the storage controller's /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image.

### Corrective Action

Recover the original file from a backup file or Snapshot(tm) copy, if possible. A copy of the file might be available from NetApp Support, although not all tape configuration files are made available there.

## Syslog Message

Option key %s is absent or badly formatted in tape configuration file %s. Default %s will be used.

## Parameters

**badOptKey** (STRING): Absent or badly formatted option keys.

**configFile** (STRING): Name of the file with the incorrect or absent option key.

**defaultOpt** (STRING): Default values used in place of the option keys.

## tape.config.fmtChkFail

### Severity

ERROR

### Description

This message occurs when a tape configuration file has a format error in a parameter that prevents use of the tape drive that the file specifies. This error can occur if a configuration file is altered, an action that is not recommended. Additional messages detail the missing or badly formatted keys. Data ONTAP® examines tape configuration files in the storage controller's /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image.

### Corrective Action

Recover the original file from a backup file or Snapshot(tm) copy, if possible. A copy of the file might be available from NetApp Support, although not all tape configuration files are made available there.

## Syslog Message

Tape configuration file %s has missing or badly formatted required keys. Format check has failed.

## Parameters

**configFile** (STRING): Name of the file with the error.

## tape.config.fmtErrKey

### Severity

ERROR

### Description

This message occurs when the file has a parameter format error in a tape configuration file that prevents the file from being used to specify tape drive parameters. Data ONTAP® examines tape configuration files in the storage controller's /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image. This error can occur if a configuration file is altered, an action that is not recommended.

### Corrective Action

Recover the original file from a backup file or Snapshot(tm) copy, if possible. A copy of the file might be available from NetApp Support, although not all tape configuration files are made available there.

## Syslog Message

Tape configuration file %s has a format error in the information appended to key %s.

## Parameters

**configFile** (STRING): Name of the file with the error.

**line** (STRING): Line containing the error.

## tape.config.infoExist

### Severity

ERROR

### Description

This message occurs when the information in a tape configuration file specifies a tape drive that Data ONTAP® already has parameters for. This error can occur if duplicate or multiple files for the same device are in the `/etc/tape_config` (7-Mode) or `/mroot/etc/tape_config` (Cluster-Mode) directory. Data ONTAP examines files in the storage controller's `tape_config` directory to dynamically qualify devices that are not part of the Data ONTAP image.

### Corrective Action

Remove any duplicates. If the device appears in the output of the "storage show tape supported" (7-Mode or Cluster-Mode nodeshell) command without the notation "(Dynamically qualified)", then Data ONTAP contains this information internally, and the file should be removed from the directory. This file might have been left over from a previous Data ONTAP release that did not have internal information for the device.

### Syslog Message

Tape device information already exists. Dynamic tape configuration file %s not used. Tape product ID already qualified.

### Parameters

**configFile** (STRING): Name of the file with the extra or duplicate information.

## tape.config.maxFilesExceeded

### Severity

ERROR

### Description

This message occurs when the maximum number of tape configuration files that Data ONTAP® will examine has been exceeded. Data ONTAP examines files in the storage controller's `/etc/tape_config` (7-Mode) or `/mroot/etc/tape_config` (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image. This error might be caused by the addition of tape configuration files retrieved from the NetApp Support site, or by user attempts to create configuration files.

### Corrective Action

Remove any files that might be duplicates, and remove any user-created files.

### Syslog Message

The maximum number of Dynamic Tape Qualification configuration files (%d) has been exceeded.

### Parameters

**maxFiles** (INT): Maximum number of dynamic qualification files allowed.

## tape.config.openFail

### Severity

ERROR

## Description

This message occurs when Data ONTAP® cannot open a file in the tape configuration directory. Data ONTAP examines files in the storage controller's /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory to dynamically qualify devices that are not part of the Data ONTAP image.

## Corrective Action

This error might be a temporary condition. If any files in the /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode) directory are opened by an editor or viewer, close them. Data ONTAP will try to open the file again for qualification when a tape operation is attempted.

## Syslog Message

Failed to open tape configuration input file %s.

## Parameters

**configFile** (STRING): File that could not be opened.

# tape.device events

## tape.device.adapterErr

### Severity

ERROR

### Description

This message occurs when an adapter error occurs that Data ONTAP® does not recognize.

### Corrective Action

If the error persists, contact NetApp technical support for assistance if the error persists.

### Syslog Message

Tape device %s: Unexpected adapter error 0x%x.

### Parameters

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**adapterError** (INT): Internal adapter error that was detected.

## tape.device.cmdAbort

### Severity

ERROR

### Description

This message occurs when a tape command fails with an error detected by the adapter. This is not an error reported by the tape drive. The adapter status values reported under this message indicate that a timeout in the transport of the command has occurred, or that other devices on the bus or loop have induced an error in the communication with this device, which might not be at fault. These errors might be of a transient nature.

### Corrective Action

Verify that the loop or bus cabling is correctly installed, and retry the operation. If the error occurs only on this device and not on others on the same interface, ensure that the external and internal cabling, if any, of



this device are securely connected. If this error occurs on more than one device, then a more general debug of the cabling infrastructure might be required. If errors of other types occur on other devices on the same bus or loop, attend to those first, because they could be the cause of this error.

### Syslog Message

Tape device %s: command aborted by host, ha\_status 0x%x.

### Parameters

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**adapterError** (INT): Internal adapter error that was detected.

## tape.device.noCompress

### Severity

ERROR

### Description

This message occurs when the tape device does not support the compression mode that Data ONTAP® attempted to set it to. This error indicates a mismatch in expectations between Data ONTAP and the tape drive over whether compression is supported. This error can occur if the the tape drive is qualified with a tape configuration file in /etc/tape\_config (7-Mode) or /mroot/etc/tape\_config (Cluster-Mode), and the file has been modified from its released condition, an action that is not recommended.

### Corrective Action

Verify whether the device documentation indicates that it supports compression. If support is indicated, then the device might need to be power-cycled to clear an abnormal condition, or it might need to be repaired. If Data ONTAP employs a tape configuration file to enable support of this drive and the file has been modified, the file should be returned to its original condition. Support using the configuration file is indicated if the device appears in the output of the "storage show tape supported" (7-Mode or Cluster-Mode nodeshell) command with the notation "(Dynamically Qualified)". Revert to a previous configuration file version, or retrieve a new version from NetApp Support, if one is available.

### Syslog Message

Tape device %s does not support compression %s.

### Parameters

**devAddr** (STRING): Device address.

**modifier** (STRING): The SCSI component that is not supported; the Mode Select Compression Page or compression capability represented in that page.

## tape.device.notSupported

### Severity

ERROR

### Description

This message occurs when Data ONTAP® does not recognize this tape drive as a supported device.

### Corrective Action

Go to [support.netapp.com/NOW/download/tools/tape\\_config/index.shtml](http://support.netapp.com/NOW/download/tools/tape_config/index.shtml) to determine whether NetApp has qualified this tape drive, and download the appropriate configuration file, if necessary. If this tape drive emulates a qualified tape drive, see [na\\_cloned\\_tapes\(5\)](#) for more information.

## Syslog Message

Data ONTAP does not recognize this tape drive: %s: %.8s%.16s%.4s

## Parameters

**address** (STRING): Address identifier of the tape device.

**vendorID** (STRING): Tape device vendor, as found in the SCSI Inquiry data field "VENDOR IDENTIFICATION".

**productID** (STRING): Product name, as found in the SCSI Inquiry data field "PRODUCT IDENTIFICATION".

**revision** (STRING): Product revision, as found in the SCSI Inquiry data field "PRODUCT REVISION LEVEL".

## tape.device.selectTimeout

### Severity

ERROR

### Description

This message occurs when the adapter cannot select the tape drive. This can be caused by the failure of the device, or by unstable loop or bus conditions that corrupt the selection signals so that the device does not detect selection request, or the adapter does not recognize the device response. This is not an error reported by the tape drive. This error might be of a transient nature.

### Corrective Action

Verify that the loop or bus cabling is correctly installed, and retry the operation. If this error occurs on more than one device, then a more general debug of the cabling infrastructure might be required. For SCSI devices, carefully examine the bus pins for damage, especially if cables have been moved since the last time this device functioned correctly. If the error persists, power-cycling the drive might clear a hang condition in the device.

## Syslog Message

Tape %s: selection timeout, ha\_status 0x%x.

## Parameters

**deviceName** (STRING): Electrical path name of the device on which the command failed.

**adapterError** (INT): Internal adapter error that was detected.

## tape.diagmsg events

### tape.diagMsg

### Severity

INFORMATIONAL

### Description

This message occurs when there is a new state of the tape diagnostic message flag. Data ONTAP® is more or less verbose in reporting SCSI check conditions depending on the flag. Using a valid tape drive address [tape], diagnostic messages can be enabled using the "mt -f [tape] diag 1" (7-Mode or Cluster-Mode nodeshell) command. You can disable the extra message using the "mt -f [tape] diag 0" command. The default value is 0.

### **Corrective Action**

There is no corrective action for this message. However, enabling diagnostic messages might cause other error messages to be generated that are usually suppressed by Data ONTAP, particularly reports of tape position. No corrective action is needed if the reports are normal. Normal reports include check conditions at the beginning or end of tape, and filemarks. It is important to understand the format of the tape data to determine whether the additional messages indicate an error.

### **Syslog Message**

Tape driver diagnostic messages %s.

### **Parameters**

**diagState** (STRING): New condition of "diag" message enable.

# tapemc events

## tapemc.alias events

### tapemc.alias.addFail

#### Severity

ERROR

#### Description

This message occurs when the automatic addition of a tape or medium changer alias fails. Data ONTAP® might attempt to add aliases automatically at boot or when devices are added, if it determines that the devices have no existing aliases.

#### Corrective Action

Check the tape and medium changer aliases to ensure that they specify the desired devices. You can view, add, or delete aliases using the "storage alias/unalias" (7-Mode or Cluster-Mode nodeshell) commands.

#### Syslog Message

Unable to retrieve and add alias for %s device %s.

#### Parameters

**deviceType** (STRING): Type of device for which an alias addition failed.

**deviceAddr** (STRING): Address of the device for which an alias addition failed.

### tapemc.alias.addOK

#### Severity

NOTICE

#### Description

This message occurs when a tape or medium changer alias is added successfully. Data ONTAP® might add aliases automatically at boot or when devices are added, if it determines that devices have no existing aliases.

#### Corrective Action

(None).

#### Syslog Message

Alias %s automatically added for %s device %s.

#### Parameters

**alias** (STRING): New alias of the device.

**deviceType** (STRING): Type of device.

**deviceAddr** (STRING): Address of the device.

### tapemc.alias.max

## Severity

ERROR

## Description

This message occurs when Data ONTAP® automatically adds a tape or media changer alias, and the new number of aliases exceeds the maximum number of aliases allowed. Automatic alias addition can occur when new devices are discovered, or when device connections are changed.

## Corrective Action

Delete any unused aliases with 'storage tape alias clear'.

## Syslog Message

Could not add alias for %s device %s. The alias count exceeds the maximum of %d aliases.

## Parameters

**deviceType** (STRING): Type of device.

**deviceAddr** (STRING): Address of the device.

**max** (INT): Maximum number of tape and media changer aliases allowed.

## tapemc.alias.threshold

### Severity

ERROR

### Description

This message occurs when Data ONTAP® automatically adds a tape or media changer alias, and the new number of aliases exceeds an internal threshold. The alias is added, but additional aliases might not be added until other aliases are deleted. Automatic alias addition can occur when new devices are discovered, or when device connections are changed.

### Corrective Action

Delete any unused aliases with 'storage tape alias clear'.

### Syslog Message

Addition of alias %s for %s device %s totals %d aliases which exceeds threshold of %d tape/media changer aliases. The maximum number of aliases allowed is %d.

### Parameters

**alias** (STRING): New alias of the device.

**deviceType** (STRING): Type of device.

**deviceAddr** (STRING): Address of the device.

**count** (INT): Current count of tape and media changer aliases.

**threshold** (INT): Internal threshold for number of tape and media changer aliases.

**max** (INT): Maximum number of tape and media changer aliases allowed.

## tapemc.device events

### tapemc.device.cmdAllocFail

**Severity**

ERROR

**Description**

This message occurs when Data ONTAP® cannot obtain a command block for sending operations to the device.

**Corrective Action**

This problem is probably a temporary resource shortage within Data ONTAP. Retry the operation.

**Syslog Message**

Failed to allocate SCSI command block for %s device %s.

**Parameters**

**deviceType** (STRING): Type of device on which the command allocation failed.

**deviceName** (STRING): Name of the device.

**tapemc.device.lunNotConn****Severity**

ERROR

**Description**

This message occurs when the tape drive, media changer, or bridge device is reported as not being connected to its dependent LUN even though it indicated that the LUN device exists. Data ONTAP® cannot initialize the dependent device.

**Corrective Action**

Consult the vendor documentation for the attached device to ensure that the device is configured correctly.

**Syslog Message**

Aborting initialization of device %s at address L%d. The attached hardware indicates that no device is available at that LUN address.

**Parameters**

**devName** (STRING): Name of the device.

**lun** (INT): LUN number of this device.

**tapemc.device.maxExceed****Severity**

ERROR

**Description**

This message occurs when more than the maximum supported number of devices of the reported type are attached to the controller.

**Corrective Action**

The reported device cannot be accessed. If other devices of the same type are attached but unused, removing them can allow the named device to be reattached and used. For FC devices, unplug the devices and wait for the messages that indicate the devices have been removed. For SCSI devices, use the

"storage adapter disable [name]" (7-Mode or Cluster-Mode nodeshell) command.

### **Syslog Message**

Maximum number of %s devices exceeded. Device %s cannot be accessed. The maximum is %d.

### **Parameters**

**deviceType** (STRING): Type of device that exceeded the maximum number allowed.

**deviceName** (STRING): Name of the device.

**deviceMax** (INT): Maximum number of devices of this device type.

## **tapemc.device.multipath**

### **Severity**

ERROR

### **Description**

This message occurs when a tape drive or media changer is reported as having a multipath connection when alias mapping is WWNN-based. WWNN-based aliasing does not support multipath connections.

### **Corrective Action**

Remove the multipath connection to the device. If multipathing is required: 1. Reboot or takeover the node and enter the boot menu (LOADER>) by pressing Ctrl+C when prompted. 2. Enter "unsetenv bootar.tape\_mc.wwnn\_alias?". 3. Enter "boot\_ontap" to reboot the node. Aliasing will now be serial number-based.

### **Syslog Message**

%s device "%s" has a multipath connection with WWNN-based alias mapping. Multipath configurations are only supported with serial number-based aliases.

### **Parameters**

**deviceType** (STRING): Type of device.

**deviceName** (STRING): Name of the device.

## **tapemc.device.noSerial**

### **Severity**

ERROR

### **Description**

This message occurs when the device can not return a serial number when requested to do so. The device can still be used, but the serial number will not be available for output.

### **Corrective Action**

This condition might be normal, especially with older device types, because SCSI-2 does not require serial numbers. If the vendor documentation indicates that the serial number should be available (look for a serial number request as an option of the Inquiry command), the device might need to be repaired.

### **Syslog Message**

Serial number not available for %s device %s, product ID %s.

## Parameters

**devType** (STRING): SCSI device type.  
**devAddr** (STRING): Device address.  
**prodID** (STRING): Vendor product ID of the device.

## tapemc.device.notExpected

### Severity

ERROR

### Description

This message occurs when Data ONTAP® expects a media changer, tape device or bridgerouter but encounters an unexpected or unknown device type.

### Corrective Action

Consult the vendor documentation for the attached device to ensure that the device is configured correctly.

### Syslog Message

Aborting device initialization for %s. Expecting a %s device but discovered an unexpected %s device type.

## Parameters

**devName** (STRING): Name of the device.  
**expectedDevType** (STRING): Expected SCSI device type.  
**actualDevType** (STRING): Actual SCSI device type discovered.

## tapemc.device.notSCSI2

### Severity

ERROR

### Description

This message occurs when Data ONTAP® requires that parallel SCSI and FC devices support, at a minimum, the SCSI-2 protocol. This device does not have the minimum support, and cannot be used with Data ONTAP unless it is upgraded.

### Corrective Action

Verify with vendor documentation that the device supports at least the SCSI-2 protocol. If it does, then the device might need to be power-cycled or repaired. Check with the vendor whether an upgrade to SCSI-2 is possible.

### Syslog Message

%s device %s not configured to support SCSI-2.

## Parameters

**devType** (STRING): SCSI device type.  
**devAddr** (STRING): Name of the device.

## tapemc.device.releaseFail



## Severity

ERROR

## Description

This message occurs when Data ONTAP® cannot remove a SCSI Persistent Reservation from this device.

## Corrective Action

Attempt the manual release of the reservation using the "storage release [tape|mc]" command (from the nodeshell in clustered environments) at the advanced privilege level. If the error persists, the device might need to be power-cycled.

## Syslog Message

Device reservation release failed on %s device %s.

## Parameters

**deviceType** (STRING): Type of device on which the reservation release failed.

**deviceName** (STRING): Name of the device.

## tapemc.device.resvConfl

### Severity

ERROR

### Description

This message occurs when Data ONTAP® detects a SCSI reservation conflict while accessing this device. Data ONTAP routinely places reservations on devices to prevent multiple simultaneous accesses by more than one storage controller. A reservation conflict indicates that an existing reservation for another host system has been detected. Conflicts might occur in some error conditions where a reservation has been left on the device inadvertently, or if a device has been moved from one system to another without a power-cycle.

### Corrective Action

If another system is using the device, then this error is normal and no corrective action should be taken to release the reservation. If the reporting system is the only user, then the reservation is likely an error.

Attempt manual release of the reservation using via the "storage release [tape|mc]" (from the nodeshell) command at the advanced privilege level. If the error persists, the device might need to be power-cycled.

### Syslog Message

Reservation conflict on %s device %s.

### Parameters

**deviceType** (STRING): Type of device on which the reservation release failed.

**deviceName** (STRING): Name of the device.

## tapemc.iopending events

### tapemc.iopending.exceeded

### Severity

INFORMATIONAL

## Description

This message occurs when Data ONTAP® detects that multiple I/Os were sent to a device that supports only one I/O at a time.

## Corrective Action

(None).

## Syslog Message

%s device %s has more than one I/O pending, but the device supports only one I/O at a time.

## Parameters

**deviceType** (STRING): Type of device that multiple I/Os were sent to.

**deviceName** (STRING): Name of the device.

# tapemc.passthru events

## tapemc.passthru.negUnderrun

### Severity

ERROR

### Description

This message occurs when a data underrun (less data is transferred than requested) happens on a read operation to the device, and the number of data bytes reported as not transferred is unexpectedly larger than the number of bytes requested. Any data involved in the transfer might not be correct. Underruns themselves can be normal if the amount of data available to transfer can be smaller than the requested amount.

### Corrective Action

Retry the operation. If the command was expected to fail, or if the failure is inconsequential, then you can ignore this condition. Otherwise, ensure that the cables and terminators (if any) are fully seated and tightened snugly. If the error persists, consider replacing the cable/terminator, host adapter or device.

### Syslog Message

%s %s inbound data transfer ended with an underrun, but the reported residue %d (bytes not transferred) exceeds the request size %d.

### Parameters

**deviceType** (STRING): Type of device.

**deviceName** (STRING): Path name of the device.

**resideBytes** (INT): Number of bytes reported as not transferred.

**requestedBytes** (INT): Number of bytes requested.

# tapemc.sas events

## tapemc.SAS.dedicatedPort

### Severity

ERROR

## Description

This message occurs when Data ONTAP® detects a SAS device that is connected to a port that is dedicated for SAS disk storage attachment, but is not a disk shelf or a bridge device.

## Corrective Action

Connect the device using another available port.

## Syslog Message

"%s" device type should not be connected to port %s, which is dedicated to disk storage attachment. Connect it using another available port.

## Parameters

**deviceType** (STRING): Type of device connected to the port.

**connectionPort** (STRING): Connection port for the device.

## tapemc.SAS.not.supported

### Severity

ERROR

### Description

This message occurs when a serial attached SCSI (SAS) tape or media changer is detected that is not allowed to be attached to this system. This is likely due to the SAS connections being reserved for disk shelf attachment. Tape drives and media changers cannot reliably share a SAS interface with disk drives.

### Corrective Action

Remove the SAS tape or media changer device.

### Syslog Message

Attaching a SAS tape or media changer device on adapter "%s" is not supported on this system.

### Parameters

**AdapterName** (STRING): Adapter name to which the tape device is attached.

# tar events

## tar.csum events

### tar.csum.csumUnsuccessful

**Severity**

NOTICE

**Description**

This message occurs when the system accesses a TAR archive and compares the calculated checksum of this archive to a stored checksum, but the checksum calculation does not succeed.

**Corrective Action**

(None).

**Syslog Message**

Checksum of %s could not be calculated.

**Parameters**

**tgz\_path** (STRING): Path to the archive for which the checksum calculation failed.

### tar.csum.extractFailed

**Severity**

ERROR

**Description**

This message occurs when the system accesses a TAR archive and compares the calculated checksum of the archive to a stored checksum, but the extraction of such an archive fails.

**Corrective Action**

(None).

**Syslog Message**

TAR extraction of %s failed with rc=%d.

**Parameters**

**tgz\_path** (STRING): Path to the archive for which extraction failed.

**rc** (INT): Error code clarifying why the extraction of the archive failed.

### tar.csum.match

**Severity**

INFORMATIONAL

**Description**

This message occurs when the system accesses a TAR archive and compares the calculated checksum of the archive to a stored checksum. In this case, the calculated and stored checksums for the archive match.

**Corrective Action**

(None).

**Syslog Message**

Stored checksum matches, not extracting %s.

**Parameters**

**tgz\_path** (STRING): Path to the archive for which the calculated and stored checksums match.

**tar.csum.mismatch****Severity**

NOTICE

**Description**

This message occurs when the system accesses a TAR archive and compares the calculated checksum of the archive to a stored checksum. In this case, the calculated and stored checksums for the TAR archive do not match.

**Corrective Action**

(None).

**Syslog Message**

Stored checksum %u does not match calculated checksum %u, extracting %s.

**Parameters**

**stored\_csum** (INT): Value of the stored checksum.

**calculated\_csum** (INT): Value of the calculated checksum.

**tgz\_path** (STRING): Path to the archive for which the calculated and stored checksums do not match.

**tar.csum.notFound****Severity**

NOTICE

**Description**

This message occurs when the system accesses a TAR archive and compares the calculated checksum of the archive to a stored checksum. In this case, the file containing the stored checksum does not exist, necessitating the extraction of the archive.

**Corrective Action**

(None).

**Syslog Message**

Stored checksum file does not exist, extracting %s.

**Parameters**

**tgz\_path** (STRING): Path to the archive for which a corresponding stored checksum file does not exist.

# tar.extract events

## tar.extract.failure

### Severity

ERROR

### Description

This message occurs during the process of extracting a TAR archive, when an error while processing a particular entry causes the extraction of the archive to fail.

### Corrective Action

(None).

### Syslog Message

Failure while extracting file (%d): '%s' %s.

### Parameters

**type** (INT): Type identifier of the TAR archive entry being processed when the error was encountered. Type identifiers and corresponding file types are as follows: 48 = Regular file, 49 = Hard link, 50 = Symbolic link, 51 = Character device, 52 = Block device, 53 = Directory, 54 = FIFO special file, 55 = Reserved.

**path** (STRING): Relative path of the TAR archive entry being processed when the error was encountered.

**errstr** (STRING): Error string providing more information about the error that caused the extraction process to fail.

## tar.extract.success

### Severity

INFORMATIONAL

### Description

This message occurs when a TAR archive has been successfully extracted.

### Corrective Action

(None).

### Syslog Message

Completed extracting %s.

### Parameters

**path** (STRING): Path to the TAR archive that has been successfully extracted.

## tar.extract.unhandledType

### Severity

ERROR

### Description

This message occurs during the process of extracting a TAR archive, when the system discovers an entry that has an unsupported type identifier.

**Corrective Action**

(None).

**Syslog Message**

File: %s has unhandled file type: %d.

**Parameters**

**path** (STRING): Relative path of this entry of the TAR archive.

**type** (INT): Type code of this entry of the TAR archive. Type identifiers and corresponding file types are as follows: 48 = Regular file, 49 = Hard link, 50 = Symbolic link, 51 = Character device, 52 = Block device, 53 = Directory, 54 = FIFO special file, 55 = Reserved.

# tcp events

## tcp.cc events

### tcp.cc.NewRenoInUpgrade

**Severity**

ERROR

**Description**

This message occurs when the system is configured to use the TCP NewReno congestion control algorithm. The TCP CUBIC algorithm is the recommended algorithm.

**Corrective Action**

The TCP NewReno congestion control algorithm is selected for the Default IPspace in 9.8 based on the system's previous configuration. Confirm that TCP NewReno is the desired algorithm. NetApp recommends the CUBIC TCP congestion control algorithm in most cases. Use the "network options congestion-control modify" command to change the congestion control algorithm.

**Syslog Message**

The system is configured to use the NewReno congestion control algorithm, which is not recommended.

**Parameters**

(None).

## tcp.change events

### tcp.change.cc.algorithm

**Severity**

ERROR

**Description**

This message occurs when an attempt to change the TCP congestion control algorithm fails.

**Corrective Action**

Check the node health, address any issues found, and then try the command again.

**Syslog Message**

An error (out of memory) occurred while changing the TCP congestion control algorithm of existing connection(s) to %s. The NewReno (RFC6582) will be used as the congestion control algorithm for those connections.

**Parameters**

**cc\_algo** (STRING): Name of the congestion control algorithm to which the change was attempted.



# tcpip events

## tcpip.conn events

### tcpip.conn.leaked

#### Severity

INFORMATIONAL

#### Description

This message occurs when a connection is leaked because it can't be closed.

#### Corrective Action

Many leaked connections could lead to degraded system performance. A power cycle will fix this.

#### Syslog Message

Connection %u to remote IP %s on vif %u using service %u over %s leaked because it could not be closed.

#### Parameters

**CID** (INT): Connection ID that was leaked.

**remoteIPAddr** (STRING): The remote IP to which the leaked connection was connected.

**VIF** (INT): The VIF on which this connection was located.

**serviceType** (INT): The service being used on the leaked connection.

**protocol** (STRING): The protocol being used on the leaked connection.

## tcpip.duplicate events

### tcpip.duplicate.ip

#### Severity

ERROR

#### Description

This message occurs when the local node detects that it and some other device are using the same IP address.

#### Corrective Action

Check the IP address configuration on both the local node and the device using the Ethernet MAC address. Modify or remove the IP address on whichever device is misconfigured.

#### Syslog Message

Duplicate IP address %s sent from link address %s.

#### Parameters

**ipAddress** (STRING): IP address under contention.

**macAddress** (STRING): Ethernet MAC address of the other device.

# telnet events

## telnet.enabled events

### telnet.enabled

**Severity**

NOTICE

**Description**

This message occurs when telnet is enabled in the cluster using the (privilege: advanced) "security protocol" command.

**Corrective Action**

(None).

**Syslog Message**

Insecure protocol service telnet has been enabled in the cluster.

**Parameters**

(None).

## telnet.firewall events

### telnet.firewall.remove

**Severity**

NOTICE

**Description**

This message occurs when telnet is automatically removed from firewall policies during a revert operation from Data ONTAP® 9.0 to 8.3. Telnet is removed because it is an insecure protocol and it is by default disabled in Data ONTAP 9.0. In Data ONTAP 8.3, the presence of telnet in firewall policies after a revert operation enables telnet. To prevent this, the telnet firewall policy is removed during the operation.

**Corrective Action**

Use the "system services firewall policy create" command to reenable telnet.

**Syslog Message**

Telnet removed from firewall policy because of Data ONTAP revert operations. Telnet is now disabled.

**Parameters**

(None).

## telnet.setsockopt events

### telnet.setsockopt.failure

**Severity**

ERROR

**Description**

This message occurs when setsockopt on telnet socket fails

**Corrective Action**

Use the error code contained in this message to determine the problem.

**Syslog Message**

Telnet setsockopt error: option %s, option value %s, error code %d (%s).

**Parameters**

**socketlevel** (STRING): This is the socket option level parameter

**socketoption** (STRING): This is the socket option value parameter

**errorcode** (INT): This is the error code returned from setsockopt

**reason** (STRING): Reason why setsockopt failed on telnet socket.

## telnet.socket events

### telnet.socket.failure

**Severity**

ERROR

**Description**

This message occurs when telnet could not create a socket.

**Corrective Action**

Change the option 'telnet.port', to try to create another socket.

**Syslog Message**

Telnet socket error: %s.

**Parameters**

**reason** (STRING): Reason why telnet could not create a socket.

### telnet.socket.timeout

**Severity**

ERROR

**Description**

This message occurs when a telnet daemon socket was not able to send data to the telnet or ssh client after multiple attempts. This message can also occur when the client window size is set to zero by the telnet or ssh client.

**Corrective Action**

Check the window size of the telnet/ssh client by typing "netstat" on the storage system. If it is zero, close the client and restart that session. Otherwise check for any network problem between telnet/ssh client and storage system and correct it.

## Syslog Message

Telnet daemon socket timed out for the %s connection.

## Parameters

**conntype** (STRING): The type of connection that timed out. The connection type can be an interactive SSH connection or a telnet connection.

# telnet.threadcreate events

## telnet.threadCreate.failure

### Severity

ERROR

### Description

This message occurs when creation of telnet daemon threads or telnet session threads fails because required virtual memory is unavailable. This happens when the system is overloaded.

### Corrective Action

Reduce the number of applications running on the storage system, then retry the operation.

### Syslog Message

Creation of %s thread failed.

### Parameters

**threadName** (STRING): Name of the thread that cannot be created.

# tls events

## tls.insecure events

### tls.insecure.cipher

#### Severity

ERROR

#### Description

This message occurs when during the TLS handshake a peer negotiates a cipher suite that is no longer considered sufficiently secure to be enabled in the default configuration. TLS connections with this peer should be be presumed insecure, and could be disrupted.

#### Corrective Action

The issue primarily needs to be resolved on the peer side. Typically, the problem is that the SSL/TLS libraries on the peer are outdated and don't support more modern cipher suites. Once all peers identified by this event have been resolved and the event is no longer being observed, it is advised to disable the cipher suite in question on this system via the "security config modify" command.

#### Syslog Message

A TLS peer with IP %s has been connected with using the insecure cipher suite %s.

#### Parameters

**ip** (STRING): IP address of the offending peer.

**cipher\_suite** (STRING): Cipher suite used by the offending peer.

### tls.insecure.protocol

#### Severity

ERROR

#### Description

This message occurs when during the TLS handshake a peer negotiates a protocol level that is no longer considered sufficiently secure to be enabled in the default configuration. TLS connections with this peer should be be presumed insecure, and could be disrupted.

#### Corrective Action

The issue primarily needs to be resolved on the peer side. Typically, the problem is that the SSL/TLS libraries on the peer are outdated and don't support more modern protocol levels. Once all peers identified by this event have been resolved and the event is no longer being observed, it is advised to disable the protocol level in question on this system via the "security config modify" command.

#### Syslog Message

A TLS peer with IP %s has been connected with using the insecure protocol level %s.

#### Parameters

**ip** (STRING): IP address of the offending peer.

**protocol\_level** (STRING): Protocol level used by the offending peer.

# tls.unused events

## tls.unused.cipher

### Severity

INFORMATIONAL

### Description

This message occurs when a cipher suite that is enabled in the configuration has not been used in any TLS handshakes for a user-specified duration of time.

### Corrective Action

Confirm that all known TLS peers have connected to the system within the specified duration of time. If this is not the case, consider increasing the used-age-threshold parameter of the "security config modify" command. Otherwise, consider disabling the unused cipher suite using the supported-cipher-suites parameter of the same command, as this will increase the security of the system with respect to TLS.

### Syslog Message

The TLS cipher suite %s has not been used for at least %s.

### Parameters

**cipher\_suite** (STRING): Cipher suite that is unused.

**duration** (STRING): Length of time (D/H/M/S) that the cipher suite has been unused.

## tls.unused.protocol

### Severity

INFORMATIONAL

### Description

This message occurs when a protocol level that is enabled in the configuration has not been used in any TLS handshakes for a user-specified duration of time.

### Corrective Action

Confirm that all known TLS peers have connected to the system within the specified duration of time. If this is not the case, consider increasing the used-age-threshold parameter of the "security config modify" command. Otherwise, consider disabling the unused protocol level using the supported-protocols parameter of the same command, as this will increase the security of the system with respect to TLS. For example, if only TLSv1.2 is in use, disabling TLSv1.1 (and lower) is advised.

### Syslog Message

The TLS protocol level %s has not been used for at least %s.

### Parameters

**protocol\_level** (STRING): Protocol level that is unused.

**duration** (STRING): Length of time (D/H/M/S) that the protocol level has been unused.

# token events

## token.copy events

### token.copy.complete.info

#### Severity

NOTICE

#### Description

This message occurs when a 'copy complete' request fails to free the reference on a remote node.

#### Corrective Action

(None).

#### Syslog Message

A 'copy complete' request on node %s and volume DSID %d failed - %s.

#### Parameters

**blade\_id** (STRING): Blade ID of the node to which the 'copy complete' request was sent.

**dsid** (LONGINT): DSID of the source volume where the token was created.

**reason** (STRING): Reason for the failure.

## token.inactive events

### token.inactive.ref.deleted

#### Severity

NOTICE

#### Description

This message occurs when a reference is deleted on a token that has existed for longer than the timeout window. The reference is deleted to clear stale entries in the cache.

#### Corrective Action

(None).

#### Syslog Message

Reference %s of token %d in volume %s was deleted.

#### Parameters

**blade\_id** (STRING): Blade ID of the token reference that was deleted.

**index** (INT): Index of the token whose reference was deleted.

**volume\_name** (STRING): Name of the volume whose token's reference was deleted.

## token.max events

## **token.max.count.in.volume**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the maximum number of tokens that can be created in a volume is determined.

### **Corrective Action**

(None).

### **Syslog Message**

Volume %s can contain up to %d tokens.

### **Parameters**

**volume\_name** (STRING): Name of the volume in which tokens can be created.

**max** (INT): Maximum number of tokens in the volume.

## **token.node events**

### **token.node.max.count.reached**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the maximum number of tokens per node is reached. Tokens cannot be created in the node unless some existing tokens expire.

#### **Corrective Action**

(None).

#### **Syslog Message**

Tokens cannot be created on this node because the maximum number of tokens per node has been reached.

#### **Parameters**

(None).

### **token.node.out.of.quorum**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a node goes out of quorum. Any references on a token by that node will be dropped.

#### **Corrective Action**

(None).



## Syslog Message

All token references from node (ID - %s) are dropped because the node went out of quorum.

## Parameters

**blade\_id** (STRING): Blade ID of the node that went out of quorum.

# token.pit events

## token.pit.file.create.info

### Severity

NOTICE

### Description

This message occurs when the token manager cannot create a PIT file in a volume. This can happen for the following reasons: insufficient space in the volume, unable to create the directory, or unable to set attributes.

### Corrective Action

(None).

### Syslog Message

PIT file creation in volume %s failed - %s.

### Parameters

**volume\_name** (STRING): Name of the volume in which PIT file creation failed.

**reason** (STRING): Reason for the failure.

## token.pit.file.deleted

### Severity

NOTICE

### Description

This message occurs when the PIT file of a volume is deleted. A new PIT file will be created automatically during the first 'token create' request for this volume.

### Corrective Action

(None).

### Syslog Message

PIT file in volume %s was deleted.

### Parameters

**volume\_name** (STRING): Name of the volume in which the PIT file was deleted.

# token.vol events

## **token.vol.cache.cleared**

### **Severity**

NOTICE

### **Description**

This message occurs when a 'volume offline' operation results in clearing the token cache for that volume.

### **Corrective Action**

(None).

### **Syslog Message**

Token cache for volume %s was cleared because the volume went offline.

### **Parameters**

**volume\_name** (STRING): Name of the volume whose token cache was cleared.

## **token.vol.is.compressed**

### **Severity**

NOTICE

### **Description**

This message occurs when a 'token create' request fails because the container is a compressed volume.

### **Corrective Action**

(None).

### **Syslog Message**

A 'token create' request in volume %s failed because the volume is compressed.

### **Parameters**

**volume\_name** (STRING): Name of the volume for which the 'token create' request failed.

## **token.vol.is.sparse**

### **Severity**

NOTICE

### **Description**

This message occurs when a 'token create' request fails because the container is a sparse volume.

### **Corrective Action**

(None).

### **Syslog Message**

A 'token create' request in volume %s failed because the volume is sparse.

### **Parameters**

**volume\_name** (STRING): Name of the volume for which the 'token create' request failed.

## **token.vol.max.count.reached**

### **Severity**

NOTICE

### **Description**

This message occurs when the maximum number of tokens per volume is reached. Tokens cannot be created in the volume unless some existing tokens expire.

### **Corrective Action**

(None).

### **Syslog Message**

Tokens cannot be created in volume %s because the maximum number of tokens per volume has been reached.

### **Parameters**

**volume\_name** (STRING): Name of the volume in which the maximum number of tokens has been reached.

# tsse events

## tsse.scan events

### tsse.scan.catalog.failed

#### Deprecated

Deprecated as of version 9.13.1.

#### Severity

ERROR

#### Description

This message occurs when inactive data compression scan related information failed to be updated on the catalog file.

#### Corrective Action

(None).

#### Syslog Message

Cold Data Scan catalog update failed on '%s' - %s.

#### Parameters

**vol** (STRING): Volume name.

**fail\_message** (STRING): Update failure reason.

### tsse.scan.catalog.updated

#### Deprecated

Deprecated as of version 9.13.1.

#### Severity

NOTICE

#### Description

This message occurs when inactive data compression scan related information updated on the catalog file.

#### Corrective Action

(None).

#### Syslog Message

Cold Data Scan catalog updated on '%s' (SE ratio local: %s, SE ratio migrate: %s, Incompressible: %llu%%, Hot data: %llu%%, Mode: %d, SE ratio cold migrate: %s).

#### Parameters

**vol** (STRING): Volume name.

**se\_ratio\_local** (STRING): SE Ratio for local reporting.

**se\_ratio\_migrate** (STRING): SE Ratio for migrate reporting.

**incomp\_data\_pct** (LONGINT): Incompressible data percentage.

**hot\_data\_pct** (LONGINT): Hot data percentage.

**scan\_mode** (INT): Scan mode.

**se\_ratio\_cold\_migrate** (STRING): SE Ratio for Cold migrate reporting.

## **tsse.scan.start.failed**

### **Severity**

NOTICE

### **Description**

This message occurs when either inactive data compression or garbage collect or garbage collect reclamation Scan fails to start.

### **Corrective Action**

(None).

### **Syslog Message**

'%s' failed to start on '%s' - %s.

### **Parameters**

**scan\_type** (STRING): Scan type.

**vol** (STRING): Volume or aggregate name.

**fail\_message** (STRING): Scan failure reason.

# ucm events

## ucm.disabled events

### ucm.disabled.slot

#### Severity

ERROR

#### Description

This message occurs when the Unified Target Adapter reads an invalid mode from flash memory. The valid supported modes are FC (0x1) or CNA (0x2). The adapter's flash memory has failed or has been corrupted. The adapter will be disabled.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Disabling adapter %s because invalid mode %d was read from nonvolatile memory.

#### Parameters

**adapter** (STRING): Name of the affected adapter.

**faultyMode** (INT): Value of the mode read from flash memory.

## ucm.mode events

### ucm.mode.autoReboot

#### Severity

INFORMATIONAL

#### Description

This message occurs when the FC adapter or converged networking adapter is reconfigured from a backup. The reconfiguration requires a reboot to enable correct adapter operation. The reboot occurs automatically.

#### Corrective Action

A change in the FC adapter or converged networking adapter mode configuration has occurred. The system is being rebooted to enable the new adapter configuration.

#### Syslog Message

Rebooting for adapters %s to become available.

#### Parameters

**adapters** (STRING): Affected adapters.

### ucm.mode.changed

#### Severity

INFORMATIONAL

## Description

This message occurs when the user changes the personality mode on a configurable adapter. The new mode will take effect on the next controller reboot.

## Corrective Action

(None).

## Syslog Message

Personality mode has changed to %s on adapter %s. Reboot the controller for the changes to take effect.

## Parameters

**mode** (STRING): Current adapter mode.

**adapter** (STRING): Name of the associated adapter.

## ucm.mode.needsReboot

### Severity

ERROR

### Description

This message occurs when the FC adapter or converged networking adapter is unable to automatically initiate a reboot after it reconfigures from a backup. The reconfiguration requires a reboot to enable correct adapter operation.

### Corrective Action

Perform a planned takeover and giveback on this controller node to enable the new adapter configuration.

### Syslog Message

The adapter %s was unable to automatically initiate a reboot after reconfiguration from the backup.

### Parameters

**adapters** (STRING): Affected adapter or adapters.

## ucm.type events

### ucm.type.changed

#### Severity

INFORMATIONAL

#### Description

This message occurs when the user changes the FC-4 type on a configurable adapter. The new type will take effect on the next controller reboot.

#### Corrective Action

(None).

#### Syslog Message

FC-4 type has changed to %s on adapter %s. Reboot the controller for the changes to take effect.

## Parameters

**type** (STRING): Current adapter type.

**adapter** (STRING): Name of the associated adapter.



# ucore events

## ucore.error events

### ucore.error

#### Severity

ERROR

#### Description

This message occurs when there is an error while processing a user space core dump.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

Could not dump core for process ID %d (%s). Reason: %s.

#### Parameters

**process\_id** (INT): Process ID of the process.  
**process\_name** (STRING): Name of the process.  
**reason** (STRING): Error condition encountered.

## ucore.limit events

### ucore.limit.exceeded

#### Severity

ERROR

#### Description

This message occurs when a userspace core is not dumped because the limit for the directory has been exceeded. This happens before the mroot is mounted, and the CompactFlash® card is used as a temporary dump location.

#### Corrective Action

Move the core off the /cfc card directory if the system has not done so automatically.

#### Syslog Message

'Exceeded limit for directory %s.'

#### Parameters

**directory\_name** (STRING): Current coredump directory.

## ucore.limits events

## ucore.limits.reached

### Severity

ERROR

### Description

This message occurs when a userspace core is not dumped because system limits for this process are reached.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Process ID %d (%s) cannot be dumped because system limits for this process are reached.

### Parameters

**process\_id** (INT): Process ID of the process.

**process\_name** (STRING): Name of the process.

## ucore.not events

### ucore.not.enabled

### Severity

ERROR

### Description

This message occurs when a userspace core cannot be dumped either because coredumps are not enabled or because the process called setuid, and coredump is disabled for processes that have changed their user credentials.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

'Process ID %d (%s) cannot be dumped because coredumps are not enabled.'

### Parameters

**process\_id** (INT): Process ID of the process.

**process\_name** (STRING): The name of the process.

## ucore.panicstring events

### ucore.panicString

### Severity

ERROR

## Description

This event occurs when a userspace application crashes, and we manage to obtain a panic string indicating the cause of the crash.

## Corrective Action

Contact customer support, and if necessary upload the core and/or the core report.

## Syslog Message

'%s'

## Parameters

**panic\_string** (STRING): The panic string of the crashed process.

**process\_name** (STRING): The name of the crashed process.

**fault\_type** (STRING): The type of fault we encountered (eg SIGSEGV, SIGHUP).

# ucore.ulimit0 events

## ucore.ulimit0

### Deprecated

This event is deprecated as of ONTAP 9.2.

### Severity

ERROR

### Description

This message occurs when a userspace core is not dumped because the core size limit is set to 0.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

'Process ID %d (%s) cannot be dumped because the coredump ulimit is set to 0.'

### Parameters

**process\_id** (INT): Process ID of the process.

**process\_name** (STRING): Name of the process.

# unowned events

## unowned.disk events

### unowned.disk.reminder

#### Severity

INFORMATIONAL

#### Description

This message occurs when the system has unowned/unassigned disks. Use the "disk assign" command to assign the disks to the systems so that they can be used.

#### Corrective Action

Use 'disk show -n' to determine which disks are unowned. Use 'disk assign' to assign ownership of the disks to a system.

#### Syslog Message

%d disks are currently unowned. Use the "disk assign" command to assign the disks to a system.

#### Parameters

**count** (INT): Number of unowned disks.

# upgrademgr events

## upgrademgr.cmd events

### upgrademgr.cmd.invoke.done

**Severity**

NOTICE

**Description**

This message occurs when the execution of the commands listed in the support file has completed.

**Corrective Action**

(None).

**Syslog Message**

The commands are invoked from the support file: %s

**Parameters**

**command** (STRING): Commands that are invoked from the support file when the automated nondisruptive upgrade (ANDU) has completed.

### upgrademgr.cmd.invoke.start

**Severity**

NOTICE

**Description**

This message occurs when the support file is invoked by an automated nondisruptive upgrade (ANDU).

**Corrective Action**

(None).

**Syslog Message**

The commands are invoked from the support file: %s

**Parameters**

**command** (STRING): Commands that are invoked from the support file when the ANDU has started.

## upgrademgr.fw events

### upgrademgr.fw.monitor.end

**Severity**

NOTICE

**Description**

This message occurs when an automated nondisruptive upgrade (ANDU) completes the firmware update monitoring job.

**Corrective Action**

(None).

**Syslog Message**

ANDU firmware update monitoring job completed. Status: %s.

**Parameters**

**status** (STRING): Status of overall firmware monitoring job.

**upgrademgr.fw.monitor.start****Severity**

NOTICE

**Description**

This message occurs when an automated nondisruptive upgrade (ANDU) starts a firmware update monitoring job.

**Corrective Action**

(None).

**Syslog Message**

ANDU firmware update monitoring job started.

**Parameters**

(None).

**upgrademgr.fw.unsupported****Severity**

NOTICE

**Description**

This message occurs when an automated nondisruptive upgrade (ANDU) runs on a platform that does not have SP/BMC API support.

**Corrective Action**

(None).

**Syslog Message**

ANDU firmware update monitoring job skipped because "%s"

**Parameters**

**reason** (STRING): ANDU firmware monitoring is not supported on this platform.

**upgrademgr.fw.update.fail****Severity**

ERROR

## Description

This message occurs when a firmware update has failed for any of the firmware components.

## Corrective Action

Contact Contact NetApp technical support. for assistance.

## Syslog Message

ANDU firmware update failed for firmware component: %s. Status: %s.

## Parameters

**firmware\_name** (STRING): Firmware component whose update has failed.

**status** (STRING): Category of failure for the firmware component whose update has failed.

# upgrademgr.node events

## upgrademgr.node.updateDone

### Severity

INFORMATIONAL

### Description

This message occurs when an update of node is completed successfully.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**node** (STRING): Node that successfully completed the update.

# upgrademgr.package events

## upgrademgr.package.imported

### Severity

INFORMATIONAL

### Description

This message occurs when a Data ONTAP® package is successfully imported into the cluster and processed by the upgrade manager's package handler.

### Corrective Action

(None).

### Syslog Message

A Data ONTAP package "%s" with version "%s" was successfully imported into the package repository.

## Parameters

**package** (STRING): Name of the imported package.  
**version** (STRING): Version of the imported package.

# upgrademgr.two events

## upgrademgr.two.stage.started

### Severity

NOTICE

### Description

This message occurs when a two-stage version update starts on the cluster where the target release is two versions later than the current version. The cluster is upgraded to intermediate version first, and then upgraded to the target version.

### Corrective Action

(None).

### Syslog Message

Started two-stage Data ONTAP upgrade to target version "%s". Intermediate version: "%s".

## Parameters

**targetversion** (STRING): Target Data ONTAP release version.  
**intermediateversion** (STRING): Intermediate Data ONTAP release version.

# upgrademgr.update events

## upgrademgr.update.cancelled

### Severity

INFORMATIONAL

### Description

This message occurs when the automated update of the cluster is canceled by the operator.

### Corrective Action

(None).

### Syslog Message

The automated update of the cluster has been canceled by the operator.

## Parameters

(None).

## upgrademgr.update.completed

### Severity

INFORMATIONAL



**Description**

This message occurs when the automated update of the cluster is complete.

**Corrective Action**

(None).

**Syslog Message**

The automated update of the cluster is complete.

**Parameters**

(None).

**upgrademgr.update.pausedErr****Severity**

ERROR

**Description**

This message occurs when the automated update of the cluster is paused as a result of an error.

**Corrective Action**

The automated update of the cluster has been paused due to the following reason: %s. Correct the cause of the error, and then issue the "cluster image resume-update" command to continue the automated update of the cluster.

**Syslog Message**

The automated update of the cluster has been paused due to the following reason: %s

**Parameters**

**pausedSummary** (STRING): Summary of the event causing the paused update.

**upgrademgr.update.pausedUser****Severity**

INFORMATIONAL

**Description**

This message occurs when the automated update of the cluster is paused as a result of an operator request.

**Corrective Action**

Issue the "cluster image resume-update" command to continue the automated update of the cluster.

**Syslog Message**

The automated update of the cluster has been paused due to a request by the operator.

**Parameters**

(None).

## **upgrademgr.update.resumed**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the automated update of the cluster is resumed from the paused state.

### **Corrective Action**

(None).

### **Syslog Message**

The automated update of the cluster has been resumed by the operator.

### **Parameters**

(None).

## **upgrademgr.update.started**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the upgrade manager begins an automated update of the cluster.

### **Corrective Action**

(None).

### **Syslog Message**

An automated update of the cluster has been started.

### **Parameters**

(None).

# usb events

## usb.device events

### usb.device.failed

#### Severity

ERROR

#### Description

This message occurs when multiple consecutive commands to the specified universal serial bus (USB) device are not completed within the allotted time. All recovery actions have been taken and the device cannot be used anymore.

#### Corrective Action

Ensure that all attached devices are supported devices and that they are seated correctly. If the problem persists, replace the attach devices. If the problem still persists, contact NetApp technical support for assistance in diagnosing a USB issue.

#### Syslog Message

Device %s attached to adapter %s has become unresponsive. Taking the device offline.

#### Parameters

**deviceName** (STRING): Name of the failed USB device.

**adapterName** (STRING): Name of the associated USB adapter.

### usb.device.initialize.failed

#### Severity

ERROR

#### Description

This message occurs when the Data ONTAP® universal serial bus (USB) adapter driver fails to initialize the device attached to the associated port in the associated adapter for one of the following reasons: Cannot set a unique address for the device; device descriptor is invalid or contains incorrect data; cannot set an active configuration for the device or the device had multiple interfaces. Note that the Data ONTAP® USB driver only supports USB 2.0 bulk-only mass storage devices.

#### Corrective Action

If the device is connected to an external USB port, try reinserting the device. If that fails, try replacing the device with a device from a different product family. If the device is connected to the mother board and the problem persists, contact NetApp technical support for assistance in diagnosing a USB issue.

#### Syslog Message

Device attached to adapter %s, in port %d failed initialization, %s.

#### Parameters

**adapterName** (STRING): Name of the associated USB adapter.

**port** (INT): Associated port number.

**error** (STRING): Reason the device initialization failed.

## usb.device.timeout

### Severity

ERROR

### Description

This message occurs when an outstanding command to the specified universal serial bus (USB) device is not completed within the allotted time. As part of the standard error handling sequence managed by the Data ONTAP® USB adapter driver, this command to the device is aborted and reissued.

### Corrective Action

Device level timeouts are a common indication of a USB link stability problem. In some cases, the link is operating normally and the specified device is having internal trouble processing I/O requests in a timely manner. In such cases, evaluate the specified device for possible replacement. Quite often the problem results from the partial failure of a component involved in the USB transport. The most common thing to check is the seating of the USB device into the USB port or the header.

### Syslog Message

Adapter %s encountered a device timeout on USB device %s.

### Parameters

**adapterName** (STRING): Name of the associated USB adapter.

**deviceName** (STRING): Name of the USB device where the timeout occurred.

## usb.device.unsupported.speed

### Severity

ERROR

### Description

This message occurs when the Data ONTAP® universal serial bus (USB) adapter driver detects a non high-speed device in the associated port.

### Corrective Action

Remove all non high speed devices attached to the system because the Data ONTAP USB adapter driver does not support non high-speed devices.

### Syslog Message

Ignoring non high-speed device detected in adapter %s on port %d.

### Parameters

**adapterName** (STRING): Name of the associated USB adapter.

**port** (INT): Port number within the adapter to which the non high-speed device is connected.

## usb.port.events

### usb.port.error

### Severity

ERROR

**Description**

This message occurs when the Data ONTAP® universal serial bus (USB) adapter driver detects an unrecoverable error on the associated port.

**Corrective Action**

If a device is attached to the associated port, try reinserting the device. If that fails, try replacing the device. If the problem persists, contact NetApp technical support for assistance in diagnosing a USB issue.

**Syslog Message**

In adapter %s, unrecoverable error detected for port %d. Ignoring attached device.

**Parameters**

**adapterName** (STRING): Name of the associated USB adapter.

**port** (INT): Associated port number.

# usbmon events

## usbmon.boot events

### usbmon.boot.device.failed

**Severity**

ERROR

**Description**

This message occurs when the Data ONTAP® module that is responsible for monitoring the health of the Universal Serial Bus (USB) boot devices determines that the associated boot device will fail all writes to the media.

**Corrective Action**

Replace the device.

**Syslog Message**

Boot device %s has become read-only.

**Parameters**

**device** (STRING): Name of the device.

### usbmon.boot.device.pfa

**Severity**

ERROR

**Description**

This message occurs when the Data ONTAP® Universal Serial Bus (USB) boot device health monitor PFA (predictive failure analysis) determines that failure is forthcoming for the associated boot device.

**Corrective Action**

Replace the device.

**Syslog Message**

Failure predicted for boot device %s.

**Parameters**

**device** (STRING): Name of the device.

## usbmon.unable events

### usbmon.unable.to.monitor

**Severity**

ERROR

**Description**

This message occurs when the Data ONTAP® module that is responsible for monitoring the health of the Universal Serial Bus (USB) boot devices cannot extract health information from the monitored device.

**Corrective Action**

Replace the device. If the problem persists, contact NetApp technical support.

**Syslog Message**

Data ONTAP could not monitor device %s.

**Parameters**

**device** (STRING): Name of the device.

# user events

## user.group events

### user.group.limit.reached

#### Severity

ERROR

#### Description

This message occurs when the number of UNIX users or groups in the cluster reaches or exceeds system limits. After the limit is reached, it will not be possible to add any new users, groups, or group members.

#### Corrective Action

Reduce the number of local users and groups stored in the cluster. Consider using directory services, such as NIS and LDAP. Alternatively, use the "vserver services limits modify" command to increase the system-wide limit.

#### Syslog Message

Total number of %s exceeds the system limit: %d.

#### Parameters

**configuration** (STRING): Name of the configuration table that is at the system limit.

**maxAllowableEntries** (INT): Maximum allowable number of local users or local group and group-members.

### user.group.reaching.limit

#### Severity

ERROR

#### Description

This message occurs when the number of UNIX users or groups in the cluster reaches various percentage thresholds of the system limit.

#### Corrective Action

Reduce the number of local users and groups stored in the cluster. Consider using directory services, such as NIS and LDAP.

#### Syslog Message

Total number of %s has reached %s of the system limit: %d.

#### Parameters

**configuration** (STRING): Name of the configuration table with the system limit.

**pct\_of\_limit** (STRING): Percentage of the configuration table's system limit reached.

**maxAllowableEntries** (INT): Maximum allowable number of local users or local group and group-members.



# useradmin events

## useradmin.added events

### useradmin.added.deleted

#### Severity

INFORMATIONAL

#### Description

A user, group, or role has been added, deleted, or modified.

#### Corrective Action

(None).

#### Syslog Message

The %s '%s' has been %s.

#### Parameters

**type** (STRING): The type of the affected useradmin name. This can be one of the following values: "user", "group", or "role".

**name** (STRING): The name of the affected user, group, or role.

**action** (STRING): The action that occurred to the subject. This can be one of the following values: "added", "deleted", or "modified".

## useradmin.lockedout events

### useradmin.lockedout.user

#### Severity

ERROR

#### Description

The user tried to log into the appliance using an invalid password and the account has been disabled. The amount of attempts a user has before being disabled is controlled by the security.passwd.lockout.numtries option.

#### Corrective Action

Another administrator must change the password of the disabled user's account.

#### Syslog Message

User '%s' is locked out of the appliance for failing authentication '%d' times.

#### Parameters

**username** (STRING): The name of the user that is now locked out from the appliance

**attempts** (INT): The number of failed attempts that caused the lockout.

# useradmin.name events

## useradmin.name.too.long

### Severity

ERROR

### Description

This message occurs when an administrative name in the cluster is longer than the maximum allowed 255 characters.

### Corrective Action

Assign administrative names that are no more than 255 characters.

### Syslog Message

%s name : %s.. is too long

### Parameters

**type** (STRING): Type of the name, which can be group, role or capability.

**invalidname** (STRING): Invalid name.

# useradmin.registry events

## useradmin.registry

### Severity

ERROR

### Description

This message occurs when the system encounters an internal registry error while reading the useradmin entries.

### Corrective Action

If the error persists, contact NetApp technical support for further assistance.

### Syslog Message

A general registry error occurred when %s.

### Parameters

**operation** (STRING): Operation in execution when the registry error occurred.

## useradmin.registry.specific

### Severity

ERROR

### Description

This message occurs when the system encounters an internal registry error while reading/writing a specific useradmin entry for a specific user.

### Corrective Action

Try to delete and re-create the user/group/role. If this does not work, then you may have to modify the registry directly.

### Syslog Message

A general registry error occurred during a %s for %s: '%s'.

### Parameters

**operation** (STRING): Operation in execution when the registry error occurred.

**type** (STRING): The type of name which is the source of the error. This can be one of the following values: "user", "group", "role", or "capability".

**name** (STRING): The name of the subject which is the source of the error.

## useradmin.rootaccess events

### useradmin.rootaccess.enabled

#### Severity

INFORMATIONAL

#### Description

The root account has been enabled or disabled because of the option security.rootaccess.enable. If this option is off, then the message appears on reboots and any time root tries to authenticate. A boot without etc/rc will not check this option, so root is temporarily enabled until the next boot.

#### Corrective Action

This event is informational-only. No action by the system administrator is necessary.

#### Syslog Message

root is %s due to a %s. Review option security.passwd.rootaccess.enable.

#### Parameters

**status** (STRING): This can be one of the following values: "enabled" or "disabled". If disabled, then root access is disabled.

**location** (STRING): Where the informational message appears. This can be in the authentication check (someone is checking root's password), reboot check (root is being disabled during a normal boot) manual change (someone changed the option manually)

# vdisk events

## vdisk.clonelogfull events

### vdisk.clonelogFull

#### Severity

NOTICE

#### Description

LUN clone creation has failed because the limit on the maximum supported entries(32705) has been exceeded.

#### Corrective Action

Retry upon the completion of the pending clone operations for this volume.

#### Syslog Message

Logging the LUN clone creation has failed because the limit on the maximum supported clone log entries in the volume %s has been exceeded.

#### Parameters

**vol\_name** (STRING): Name of the volume for which the LUN clone creation has failed.

# vdom events

## vdom.aggr events

### vdom.aggr.move.internal

#### Severity

INFORMATIONAL

#### Description

This message occurs when an individual mounted volume becomes unavailable through the source node after an aggregate relocation or giveback event.

#### Corrective Action

The destination node receives a notification that the volume has been mounted locally and updates the cluster accordingly.

#### Syslog Message

Availability of a volume in aggregate %s has been lost to the source node temporarily following aggregate relocation or giveback.

#### Parameters

**aggr** (STRING): Name of the aggregate that contains the volume.

## vdom.unmapped events

### VDOM.unmapped.event

#### Severity

NOTICE

#### Description

This message occurs for Virtual Disk Object Manager (VDOM) errors that do not currently map to specific VDOM EMS events.

#### Corrective Action

(None).

#### Syslog Message

VDOM error code %u occurred without a specific EMS event mapping, file: %s, line %u.

#### Parameters

**vdom\_error** (INT): VDOM error code.

**filename** (STRING): Name of file containing error.

**linenum** (INT): Line number in file containing error.

# vf events

## vf.cleanup events

### vf.cleanup.error

#### Severity

ERROR

#### Description

This event occurs when some subsystems are unable to cleanup properly when a local host vfiler is destroyed or when a partner vfiler is being shutdown during giveback.

#### Corrective Action

Log a support case with NetApp technical support. Try starting the vfiler, stopping and destroying it again.

#### Syslog Message

vfiler %s: %s subsystem could not cleanup properly.

#### Parameters

**vfName** (STRING): The vfiler being shutdown or destroyed.

**module** (STRING): The module which could not cleanup vfiler specific information properly.

# vfiler events

## vfiler.unable events

### vFiler.unable.initModules

#### Severity

ERROR

#### Description

This message occurs during vFiler(tm) unit discovery, if one of the vFiler(tm) unit module initialization functions fails.

#### Corrective Action

(None).

#### Syslog Message

Could not initialize modules for vFiler unit "%s".

#### Parameters

**vfName** (STRING): The vFiler unit on the source.

# vfilermigrate events

## vfilermigrate.diskownershipmismatch events

### vfilermigrate.diskOwnershipMismatch

**Severity**

ERROR

**Description**

The SnapMover feature is licensed but the disk ownership model is not using the software-based disk ownership model.

**Corrective Action**

Boot the filer into maintenance mode and use the 'disk assign' or 'disk upgrade\_ownership' command to assign disk ownership.

**Syslog Message**

SnapMover is licensed but the system is not using software-based disk ownership.

**Parameters**

(None).



# vifmgr events

## vifmgr.bcastdomain events

### vifmgr.bcastdomain.portadd

**Severity**

NOTICE

**Description**

This message occurs when a network port is added to a broadcast domain.

**Corrective Action**

(None).

**Syslog Message**

Network port %s on node %s was added to broadcast domain %s in IPspace %s.

**Parameters**

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**broadcastDomain** (STRING): Name of the broadcast domain associated with the port.

**ipspace** (STRING): Name of the IPspace associated with the port.

### vifmgr.bcastdomain.portrem

**Severity**

NOTICE

**Description**

This message occurs when a network port is removed from a broadcast domain.

**Corrective Action**

(None).

**Syslog Message**

Network port %s on node %s was removed from broadcast domain %s in IPspace %s.

**Parameters**

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**broadcastDomain** (STRING): Name of the broadcast domain associated with the port.

**ipspace** (STRING): Name of the IPspace associated with the port.

## vifmgr.bd events

### vifmgr.bd.badmtu

## Severity

ERROR

## Description

This message occurs when an MTU is entered on a Broadcast Domain that its member ports cannot support. The ports that failed will continue using their present values.

## Corrective Action

Change the MTU of the Broadcast Domain to a value supported by all member ports.

## Syslog Message

In IPspace %s, Broadcast Domain %s has an MTU of %d, which is not supported by the following ports: %s.

## Parameters

**ipspace** (STRING): Name of the IPspace in which the Broadcast Domain resides.

**broadcast\_domain** (STRING): Name of the Broadcast Domain.

**mtu** (INT): MTU of the Broadcast Domain.

**port\_names** (STRING): Ports in the Broadcast Domain that are failing.

## vifmgr.bd.conflicts

### Severity

ERROR

### Description

This message occurs when the administrator modifies properties of a node management LIF while the node is out of quorum, and the broadcast domain of the home port, current port, and subnet are not all the same.

### Corrective Action

The node management LIF's configuration before the node went out of quorum was automatically applied. You can change this configuration now that the node is online.

### Syslog Message

LIF %s had mismatching broadcast domains. The previous values were: home-port (%s), current-port (%s), and subnet (%s).

### Parameters

**LIF** (STRING): Name of the LIF.

**home\_port\_bd** (STRING): Broadcast domain name of this LIF's home port prior to going out of quorum.

**current\_port\_bd** (STRING): Broadcast domain name of this LIF's current port prior to going out of quorum.

**sub\_bd** (STRING): Broadcast domain name of this LIF's subnet prior to going out of quorum.

## vifmgr.bgp events

### vifmgr.bgp.connected

### Severity

NOTICE

**Description**

This message occurs when Logical Interface Manager (VifMgr) connects to BGP daemon.

**Corrective Action**

(None).

**Syslog Message**

VifMgr connected to BGP daemon on the node.

**Parameters**

(None).

**vifmgr.bgp.defaults.upgrade****Severity**

NOTICE

**Description**

This message occurs when the Autonomous System Number (ASN) in "bgp defaults" changes from the reserved number to the default number (65501).

**Corrective Action**

(None).

**Syslog Message**

Autonomous System Number (ASN) in "bgp defaults" was changed from the reserved number (%d) to the default number (65501).

**Parameters**

**old\_asn** (INT): Reserved Autonomous System Number (ASN) in "bgp defaults" before it is changed.

**vifmgr.bgp.disconnected****Severity**

ERROR

**Description**

This message occurs when the connection between the Logical Interface Manager (VifMgr) and the BGP daemon is broken. The connection is automatically reestablished. In case the connection is not reestablished, any Virtual IP LIFs might fail over to another node.

**Corrective Action**

Verify that the BGP daemon restarted by looking for the "vifmgr.bgp.connected" event. If the BGP daemon fails to restart, or if the condition persists, contact NetApp technical support.

**Syslog Message**

The connection between VifMgr and the BGP daemon on the node is broken.

**Parameters**

(None).

## vifmgr.bgp.peerGroupDown

### Severity

NOTICE

### Description

This message occurs when a Border Gateway Protocol (BGP) peer group goes down on a given IPspace. If this is the last active BGP session on this IPspace, vifmgr.bgp.vserverDown message gets generated and any virtual IP (VIP) LIF or LIFs in the data Vserver or data Vservers of this IPspace cannot be hosted on this node without an active BGP session.

### Corrective Action

(None).

### Syslog Message

The BGP peer group %s on IPspace %s is down.

### Parameters

**peer\_group** (STRING): Name of the BGP peer group.

**peer\_ipspace** (STRING): IPspace of the BGP peer group.

## vifmgr.bgp.peerGroupUp

### Severity

NOTICE

### Description

This message occurs when a BGP peer group goes up on a given IPspace. If this is the first active BGP session on this IPspace, vifmgr.bgp.vserverUp message gets generated and any VIP LIF or LIFs in the data Vserver or data Vservers of this IPspace can now be hosted on this node.

### Corrective Action

(None).

### Syslog Message

The BGP peer group %s on IPspace %s is up.

### Parameters

**peer\_group** (STRING): Name of the BGP peer group.

**peer\_ipspace** (STRING): IPspace of the BGP peer group.

## vifmgr.bgp.vserverDown

### Severity

NOTICE

### Description

This message occurs when a BGP down event is received for a system Vserver. Any VIP LIFs in the associated data Vservers can no longer handle network traffic.

### Corrective Action

(None).

### Syslog Message

VifMgr received a Vserver BGP down event for Vserver %s in IPspace %s.

### Parameters

**vserver** (STRING): Name of the Vserver.

**ipspace** (STRING): Name of the IPspace where the Vserver resides.

## vifmgr.bgp.vserverUp

### Severity

NOTICE

### Description

This message occurs when a BGP up event is received for a system Vserver. Any VIP LIFs in the associated data Vservers can now handle network traffic.

### Corrective Action

(None).

### Syslog Message

VifMgr received a Vserver BGP up event for Vserver %s in IPspace %s.

### Parameters

**vserver** (STRING): Name of the Vserver.

**ipspace** (STRING): Name of the IPspace where the Vserver resides.

## vifmgr.cdb events

### vifmgr.cdb.merge.dupip

### Severity

ALERT

### Description

This message occurs when attempting to merge a LIF from the configuration database (CDB) to the replicated database (RDB), and the LIF from the CDB is found to have the same address as an existing LIF in the RDB.

### Corrective Action

Modify the IP address of the conflicting LIF configured in the database.

### Syslog Message

IP address %s is in use by multiple LIFs: local LIF %s (on node %s) and LIF %s (in IPspace %s) configured in the cluster database.

### Parameters

**vifip** (STRING): IP address of the affected LIF.

**locallif** (STRING): Name of the local LIF.

**node** (STRING): Name of the node on which the local LIF resides.  
**dblif** (STRING): Name of the conflicting LIF configured in the database.  
**ipspace** (STRING): Name of the IPspace.

## **vifmgr.cdb.unavail**

### **Severity**

NOTICE

### **Description**

This message occurs when the configuration database (CDB) for a node is not accessible. The node will continue without the CDB.

### **Corrective Action**

(None).

### **Syslog Message**

The CDB for node %s is not accessible.

### **Parameters**

**node** (STRING): Name of the node.

## **vifmgr.clntdev events**

### **vifmgr.clntdev.removed**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a remote device is removed from the cluster or management network.

#### **Corrective Action**

(None).

#### **Syslog Message**

Remote device %s (model: %s) was removed from the %s network.

#### **Parameters**

**device** (STRING): Device name of the connected switch or node.  
**model** (STRING): Model name of the remote device.  
**network** (STRING): Network where the device was detected.

### **vifmgr.clntdev.updated**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a new remote device is detected on the cluster or management network.

## Corrective Action

(None).

## Syslog Message

Remote device %s (model: %s, version: %s, IP address: %s) has been detected on the %s network.

## Parameters

**device** (STRING): Device name of the connected switch or node.

**model** (STRING): Model name of the remote device.

**version** (STRING): Software version of the remote device.

**address** (STRING): IP address of the remote device.

**network** (STRING): Network where the device was detected.

# vifmgr.clus events

## vifmgr.clus.linkdown

### Severity

EMERGENCY

### Description

This message occurs when a cluster link goes down unexpectedly.

### Corrective Action

Verify that the cable is plugged in properly to both the NIC and the switch, and that the switch is powered on. If these conditions are met and the cluster link remains down, search the knowledgebase of the NetApp support web site for the "cluster network degraded" keyword.

### Syslog Message

The cluster port %s on node %s has gone down unexpectedly.

### Parameters

**portname** (STRING): Name of the affected port.

**nodename** (STRING): Name of the node on which the affected port resides.

# vifmgr.cluscheck events

## vifmgr.cluscheck.assigned

### Severity

NOTICE

### Description

This message occurs when the named cluster logical interface (LIF) was not assigned to any network port and is now assigned to a port.

### Corrective Action

(None).

## Syslog Message

Cluster LIF %s (node %s) is now assigned to port %s.

## Parameters

**lif\_name** (STRING): Name of cluster LIF that was not assigned to a network port but now is.

**node\_name** (STRING): Name of the node on which the cluster LIF that was not assigned to a port now resides.

**port\_name** (STRING): Name of the port to which the LIF is now assigned.

## vifmgr.cluscheck.badping

### Severity

NOTICE

### Description

This message occurs when the ping test for cluster logical interface (LIF) interconnectivity displays unexpected results.

### Corrective Action

Ensure connectivity among the cluster LIFs. Several issues that can cause this problem. Some things to check include the following: - Port connectivity to the correct network or switch - Link connectivity - Proper switch configuration - Proper LIF configuration

## Syslog Message

Cluster check ping command (%s) failed: %d, %s.

## Parameters

**cmd\_line** (STRING): Ping test command line.

**exit\_status** (INT): Ping test result.

**cmd\_output** (STRING): Ping test command output.

## vifmgr.cluscheck.crcerrors

### Deprecated

Deprecated as of version 9.8.0 and replaced by vifmgr.cluscheck.hwerrors.

### Severity

ALERT

### Description

This message occurs when a network device reports a high number of observed hardware errors, such as CRC errors, length errors, alignment errors, or dropped frames.

### Corrective Action

The errors could be originating from the specified port, a remote port, or a port on another component of the network. Check the statistics for both the port and the switch. Contact NetApp technical support for assistance and specific instructions.

## Syslog Message

Port %s on node %s is reporting a high number of observed hardware errors, possibly CRC errors.



## Parameters

**portname** (STRING): Name of the device that is reporting a high number of hardware errors.

**nodename** (STRING): Name of the node on which the port resides.

## vifmgr.cluscheck.ctdpktloss

### Severity

ALERT

### Description

This message occurs when continued packet loss during the ping test is observed. Continued packet loss means multiple failures of the ping test persisted over several iterations of the test. The ping test tests the connectivity of a cluster logical interface (LIF) to all other cluster logical interfaces (LIFs).

### Corrective Action

Ensure connectivity of the cluster LIFs. There are a number of issues that can cause this problem. Some things to check: - is the port connected to the correct network/switch - link light - switch configuration

### Syslog Message

Continued packet loss when pinging from cluster lif %s (node %s) to cluster lif %s (node %s).

## Parameters

**src\_vif** (STRING): Name of the source cluster LIF.

**src\_node** (STRING): Node name of the source cluster LIF.

**dst\_vif** (STRING): Name of the destination cluster LIF.

**dst\_node** (STRING): Node name of the destination cluster LIF.

## vifmgr.cluscheck.droppedall

### Severity

ALERT

### Description

This message occurs when all packets of the ping test are dropped. The ping test tests the connectivity of a cluster logical interface (LIF) to all other cluster logical interfaces (LIFs).

### Corrective Action

Ensure connectivity of the cluster LIFs. There are a number of issues that can cause this problem. Some things to check: - is the port connected to the correct network/switch - link light - switch configuration

### Syslog Message

Total packet loss when pinging from cluster lif %s (node %s) to cluster lif %s (node %s).

## Parameters

**src\_vif** (STRING): Name of the source cluster LIF.

**src\_node** (STRING): Node name of the source cluster LIF.

**dst\_vif** (STRING): Name of the destination cluster LIF.

**dst\_node** (STRING): Node name of the destination cluster LIF.

## vifmgr.cluscheck.droppedlarge

### Severity

ALERT

### Description

This message occurs when all of the packets in the large packet ping test are dropped. This most likely suggests an MTU misconfiguration. The ping test tests the connectivity of a cluster logical interface (LIF) to all other cluster logical interfaces (LIFs).

### Corrective Action

First, resolve any outstanding instances of vifmgr.cluscheck.mtumismatch. Second, verify that the cluster switches themselves are configured to support the same MTU as configured for the cluster LIFs.

### Syslog Message

Partial packet loss when pinging from cluster lif %s (node %s) to cluster lif %s (node %s).

### Parameters

**src\_vif** (STRING): Name of the source cluster LIF.  
**src\_node** (STRING): Node name of the source cluster LIF.  
**dst\_vif** (STRING): Name of the destination cluster LIF.  
**dst\_node** (STRING): Node name of the destination cluster LIF.

## vifmgr.cluscheck.droppednone

### Severity

NOTICE

### Description

This message occurs when where was some ping test failure that is now no longer occurring.

### Corrective Action

(None).

### Syslog Message

No packet loss when pinging from cluster lif %s (node %s) to cluster lif %s (node %s).

### Parameters

**src\_vif** (STRING): Name of the source cluster LIF.  
**src\_node** (STRING): Node name of the source cluster LIF.  
**dst\_vif** (STRING): Name of the destination cluster LIF.  
**dst\_node** (STRING): Node name of the destination cluster LIF.

## vifmgr.cluscheck.hwerrors

### Severity

ALERT

### Description

This message occurs when a network device reports a high number of observed hardware errors, such as CRC errors, length errors, alignment errors, or dropped frames.

## Corrective Action

The errors could be originating from the specified port, a remote port, or a port on another component of the network. Check the statistics for both the port and the switch. Contact NetApp technical support for assistance and specific instructions.

## Syslog Message

Port %s on node %s is reporting a high number (at least 1 per %d packets) of observed hardware errors (CRC, length, alignment, dropped).

## Parameters

**portname** (STRING): Name of the device that is reporting a high number of hardware errors.

**nodename** (STRING): Name of the node on which the port resides.

**threshold** (INT): Threshold at which this message is logged, at least 1 error per threshold.

## vifmgr.cluscheck.l2ping

### Severity

EMERGENCY

### Description

This message occurs when the l2ping test fails. This test verifies L2 reachability between cluster ports. This condition typically results in loss of communication between the two nodes.

### Corrective Action

Verify that the cable is plugged in properly to both the NIC and the switch, and that the switch is powered on. If these conditions are met and the cluster link remains down, contact NetApp technical support to investigate further.

### Syslog Message

l2ping failure when pinging from cluster port %s (node %s) to cluster port %s (node %s).

### Parameters

**src\_prt** (STRING): Name of the source cluster port.

**src\_node** (STRING): Node name of the source cluster port.

**dst\_port** (STRING): Name of the destination cluster port.

**dst\_node** (STRING): Node name of the destination cluster port.

## vifmgr.cluscheck.lifdown

### Severity

NOTICE

### Description

This message occurs when the named cluster logical interface (LIF) is configured administratively "down".

### Corrective Action

Use the "network interface modify" command to change the status-admin value of the LIF to "up".

### Syslog Message

Cluster LIF %s is administratively "down".

## Parameters

**lif\_name** (STRING): Name of the LIF that is currently inactive.

## vifmgr.cluscheck.lifup

### Severity

NOTICE

### Description

This message occurs when a cluster logical interface (LIF) was configured administratively "down" and is now configured administratively "up".

### Corrective Action

(None).

### Syslog Message

Cluster LIF %s is administratively "up".

## Parameters

**lif\_name** (STRING): Name of the LIF that is currently active.

## vifmgr.cluscheck.mtumatch

### Severity

ALERT

### Description

This message occurs when the Maximum Transmission Unit (MTU) of a port assigned to a cluster logical interface (LIF) did not match the MTU of a port assigned to a different cluster logical interface but now does match.

### Corrective Action

(None).

### Syslog Message

MTU %s on cluster port %s (node %s) now matches MTU %s on cluster port %s (node %s).

## Parameters

**src\_mtu** (STRING): MTU of the source cluster LIF port.

**src\_port** (STRING): Name of the source cluster LIF port.

**src\_node** (STRING): Name of the source node on which source cluster LIF port resides.

**dst\_mtu** (STRING): MTU of the destination cluster LIF port.

**dst\_port** (STRING): Name of the destination cluster LIF port.

**dst\_node** (STRING): Name of the node on which destination cluster LIF port resides.

## vifmgr.cluscheck.mtumismatch

### Severity

ALERT

## Description

This message occurs when the Maximum Transmission Unit (MTU) of a port assigned to a cluster logical interface (LIF) does not match the MTU of a port assigned to a different cluster logical interface. This condition may cause loss of packets between these two LIFs.

## Corrective Action

Ensure that the MTUs of the LIFs are identical. The "network port show" command will display the port MTU settings. Any mismatched MTUs can be modified from the "network port" subcommand.

## Syslog Message

MTU %s on cluster port %s (node %s) does not match MTU %s on cluster port %s (node %s).

## Parameters

**src\_mtu** (STRING): MTU of the source cluster LIF port.  
**src\_port** (STRING): Name of the source cluster LIF port.  
**src\_node** (STRING): Name of the source node on which source cluster LIF port resides.  
**dst\_mtu** (STRING): MTU of the destination cluster LIF port.  
**dst\_port** (STRING): Name of the destination cluster LIF port.  
**dst\_node** (STRING): Name of the node on which destination cluster LIF port resides.

## vifmgr.cluscheck.notassigned

### Severity

ALERT

### Description

This message occurs when the named cluster logical interface (LIF) is not assigned to any network port.

### Corrective Action

Verify that the LIF's current port is configured administratively "up". Use the "network interface show" command to determine the current port of the LIF. Use the "network port show -fields up-admin" command to show the port's administrative state. Ensure that the network cabling connecting the remote network port is in place and secure. Ensure that the remote network port is configured administratively "up". If the port's link state is "up", attempt to migrate the cluster LIF to another cluster port. Use the "network port show" command to view the port's link state. If the current port is "up", use the "network interface migrate" command to migrate the LIF to another port.

### Syslog Message

Cluster LIF %s (node %s) is not assigned to any port.

### Parameters

**lif\_name** (STRING): Name of the cluster LIF.  
**node\_name** (STRING): Name of the node on which the cluster LIF resides.

## vifmgr.cluscheck.sameRmtDev

### Severity

ERROR

### Description

This message occurs when the home ports of the two cluster LIFs used for internal cluster replicated database (RDB) operations are connected to the same cluster switch. These LIFs should be connected to

different cluster switches for optimal cluster resiliency.

### Corrective Action

Use the "network port show -ipSpace Cluster -fields remote-device-id" command to verify that the home ports of the specified cluster LIFs are connected to different cluster switches. If necessary, re-home one of the cluster LIFs to a port that is connected to the other switch. For help, search the NetApp knowledge base for "ONTAP Cluster interconnect management communication is improperly configured".

### Syslog Message

Cluster LIFs "%s" and "%s" on node "%s" are used for RDB operations. The home ports of these LIFs are connected to the same cluster switch: "%s".

### Parameters

**clus1\_lif** (STRING): Name of a cluster LIF.  
**clus2\_lif** (STRING): Name of a cluster LIF.  
**node** (STRING): Node name of the cluster LIFs.  
**remote\_device** (STRING): Name of the cluster switch.

## vifmgr.clusterbd events

### vifmgr.clusterBD.Partition

#### Severity

ERROR

#### Description

This message occurs when ONTAP detects that the cluster broadcast domain is partitioned; some local ports in the cluster broadcast domain do not have layer-2 connectivity to one another. Cluster LIFs could be at risk of connectivity issues.

#### Corrective Action

Check the network cabling of the ports in this broadcast domain to the network switch. Check the broadcast domain configuration on the storage system and the network port configuration on the switch for any misconfigurations. Run the "network port reachability show -detail" command to view which ports are unreachable. To view information about broadcast domains and ports on the cluster, run the "network port broadcast-domain show -instance" and "network port show -instance" commands.

#### Syslog Message

The cluster broadcast domain is partitioned on node "%s". LIFs hosted on the ports in this broadcast domain might be at risk of seeing connectivity issues.

#### Parameters

**nodeName** (STRING): Name of the node for which the local ports in the given broadcast domain are partitioned.

## vifmgr.clusterlfoconf events

### vifmgr.clusterLFOConf

**Severity**

NOTICE

**Description**

This message occurs when the failover policy and failover group for a cluster LIF are restored to their default values. It is necessary to restore the default values for these settings to ensure non-disruptive operation.

**Corrective Action**

(None).

**Syslog Message**

Setting failover settings for cluster LIF %s to avoid disruption.

**Parameters**

LIF (STRING): Name of the Cluster LIF that was modified.

## vifmgr.dbase events

### vifmgr.dbase.checkerror

**Severity**

ALERT

**Description**

This message occurs when the logical interface manager (VIFMgr) experiences an error while verifying cluster database consistency. Some configured objects, such as logical interfaces (LIFs) might not be hosted properly as a result.

**Corrective Action**

Contact NetApp technical support and reference the "vifmgr.dbase.checkerror" error.

**Syslog Message**

VIFMgr experienced an error verifying cluster database consistency. Some LIFs might not be hosted properly as a result.

**Parameters**

(None).

## vifmgr.displaced events

### vifmgr.displaced.lifHomed

**Severity**

NOTICE

**Description**

This message occurs when the home port of a displaced logical interface (LIF) has been restored to its original home port.

**Corrective Action**

(None).

**Syslog Message**

Home port of displaced LIF %s (Vserver: %s, IP: %s) restored to node %s, port %s.

**Parameters**

**vifName** (STRING): Name of the LIF on which the home port is being updated.

**vsrvId** (STRING): ID of the Vserver on which the LIF resides.

**vifIp** (STRING): IP address of the LIF on which the home port is being updated.

**toNode** (STRING): Name of the node to which the LIF's home port is being updated.

**toPort** (STRING): Name of the port to which the LIF's home port is being updated.

## vifmgr.extroute events

### vifmgr.extRoute.addTable

**Severity**

NOTICE

**Description**

This message occurs when a new external route table is added successfully.

**Corrective Action**

(None).

**Syslog Message**

External route table %s was added successfully.

**Parameters**

**table** (STRING): Route table name.

### vifmgr.extRoute.ipCreate

**Severity**

NOTICE

**Description**

This message occurs when a LIF was successfully added to or updated in external route tables.

**Corrective Action**

(None).

**Syslog Message**

LIF with IP address %s is mapped to external port ID %s in external route tables.

**Parameters**

**vifip** (STRING): IP address of the logical interface.

**extportid** (STRING): External port ID on which this LIF was created.



## **vifmgr.extRoute.ipDelete**

### **Severity**

NOTICE

### **Description**

This message occurs when a LIF was successfully deleted from external route tables.

### **Corrective Action**

(None).

### **Syslog Message**

LIF with IP address %s was removed from external route tables.

### **Parameters**

**vifip** (STRING): IP address of the logical interface.

## **vifmgr.extRoute.ipMove**

### **Severity**

NOTICE

### **Description**

This message occurs when a LIF was successfully moved to a new external port ID in external route tables.

### **Corrective Action**

(None).

### **Syslog Message**

LIF with IP address %s was moved to external port ID %s in external route tables.

### **Parameters**

**vifip** (STRING): IP address of the logical interface.

**extportid** (STRING): External port ID to which this LIF was moved.

## **vifmgr.extRoute.noAddTable**

### **Severity**

ERROR

### **Description**

This message occurs when ONTAP fails to add a new external route table to the cloud network. As a result, IP packets routed to the subnet associated with the external route table will not be possible.

### **Corrective Action**

1. Add the external route table manually using the tools provided by the cloud provider.
2. Check for cloud authorization and resource permission issues.

### **Syslog Message**

Failed to add external route table %s (error %s).

## Parameters

**table** (STRING): Route table name.  
**error** (STRING): Error that is reported.

## vifmgr.extRoute.noConn

### Severity

ERROR

### Description

This message occurs when the system cannot verify connectivity to the external mediator server.

### Corrective Action

Use the "network route show -vserver Cluster" command to verify that a suitable route exists to the mediator. Ensure that the external network security configuration enables communication between all nodes and the mediator.

### Syslog Message

Failed to verify connectivity with the mediator server, IP %s (error %s).

## Parameters

**mediatorip** (STRING): IP address of the mediator.  
**error** (STRING): Error that is reported.

## vifmgr.extRoute.nolpCreate

### Severity

ERROR

### Description

This message occurs when a LIF cannot be added to or updated in external route tables.

### Corrective Action

If an entry is found in the external route tables for the given IP address and external port ID, remove that entry from the tables using the tools provided by the cloud provider. Note that any subsequent LIF created with the same IP address will replace any entry remaining in the external route table.

### Syslog Message

Failed to successfully map LIF with IP address %s to external port ID %s in external route tables (error %s).

## Parameters

**vifip** (STRING): IP address of the logical interface.  
**extportid** (STRING): External port ID on which this LIF was to be created.  
**error** (STRING): Error that is reported.

## vifmgr.extRoute.nolpDelete

### Severity

ERROR

## Description

This message occurs when ONTAP fails to update the external route tables maintained by the cloud provider to reflect a LIF removal. As a result, IP packets to the IP address might still be routed to the deleted external port ID.

## Corrective Action

1. Remove the entry from the tables using the tools provided by the cloud provider if an entry is found in the external route tables for the given IP address and external port ID. 2. Verify connectivity, authentication, and permissions to the cloud endpoint using the tools provided by the cloud provider.

## Syslog Message

Failed to remove LIF with IP address %s from external route tables (error %s).

## Parameters

**vifip** (STRING): IP address of the logical interface.

**error** (STRING): Error that is reported.

## vifmgr.extRoute.noIpMove

### Severity

ERROR

### Description

This message occurs when ONTAP fails to update the external route tables maintained by the cloud provider to reflect a LIF move. As a result, IP packets to the IP address might still be routed to the previously assigned external port ID.

### Corrective Action

1. Remove the entry from the tables using the tools provided by the cloud provider if an entry is found in the external route tables for the given IP address and external port ID. 2. Verify connectivity, authentication and permissions to the cloud endpoint using the tools provided by the cloud provider. 3. Admin the status of the LIF associated with the IP down and up to force the IP movement. "network interface modify -vserver [x] -lif [xx] -status-admin down" "network interface modify -vserver [x] -lif [xx] -status-admin up"

### Syslog Message

Failed to move LIF. Error: %s

### Parameters

**error** (STRING): Error that is reported.

## vifmgr.extRoute.noRmTable

### Severity

NOTICE

### Description

This message occurs when ONTAP fails to remove the external route table from the cloud network. As a result, IP packets routed to the subnet associated with the external route table might still be valid in the virtual private cloud.

## Corrective Action

1. Remove the external route table manually using the tools provided by the cloud provider.
2. Check for cloud authorization and resource permission issues.

## Syslog Message

Failed to remove external route table %s (error %s).

## Parameters

**table** (STRING): Route table name.

**error** (STRING): Error that is reported.

## vifmgr.extRoute.rmTable

### Severity

NOTICE

### Description

This message occurs when an external route table is removed.

### Corrective Action

(None).

## Syslog Message

External route table %s was removed successfully.

## Parameters

**table** (STRING): Route table name.

## vifmgr.gcp events

### vifmgr.gcp.NlbdDown

### Severity

ERROR

### Description

This message occurs when the network load balancing listener becomes unresponsive. This event causes data LIFs to be moved to a healthier node.

### Corrective Action

Reboot the node.

## Syslog Message

The network load balancing listener is unresponsive.

## Parameters

(None).

## vifmgr.gcp.NlbdUp

### Severity

NOTICE

### Description

This message occurs when the network load balancing listener recovers after a period of unresponsiveness. If data LIFs assigned to this node have been configured to automatically revert to their home port, this will cause them to return to their home port.

### Corrective Action

(None).

### Syslog Message

The network load balancing listener has recovered after a period of unresponsiveness.

### Parameters

(None).

## vifmgr.hm events

### vifmgr.hm.promoted

### Severity

NOTICE

### Description

This message occurs when a health monitor has been fully activated for an IPspace. A fully activated health monitor marks ports as unhealthy if the port fails the monitor's health check.

### Corrective Action

(None).

### Syslog Message

The %s has been promoted on IPspace %s.

### Parameters

**monitor** (STRING): Health monitor name.

**ipspace** (STRING): IPspace name for which the health monitor has been promoted.

## vifmgr.hv events

### vifmgr.hv.nohostconfig

### Severity

ALERT

### Description

This message occurs when no Hypervisor host configuration is available. This is required by ONTAP® software to monitor the health of the hypervisor uplinks.

### Corrective Action

Verify the Hypervisor host configuration. Without it, ONTAP cannot monitor the health of the hypervisor uplinks, preventing data LIFs from failing over if the uplink were to go down.

### Syslog Message

The %s health monitor cannot access the Hypervisor host configuration.

### Parameters

**monitor** (STRING): Health monitor name.

## vifmgr.hv.nosnmpconfig

### Severity

ALERT

### Description

This message occurs when no Hypervisor SNMP configuration is available. This is required by ONTAP® software to monitor the health of the hypervisor uplinks.

### Corrective Action

Verify the Hypervisor SNMP configuration. Without it, ONTAP cannot monitor the health of the hypervisor uplinks, preventing data LIFs from failing over if the uplink were to go down.

### Syslog Message

The %s health monitor cannot access the required Hypervisor SNMP configuration.

### Parameters

**monitor** (STRING): Health monitor name.

## vifmgr.hv.nouplinkmapping

### Severity

ALERT

### Description

This message occurs when no Hypervisor uplink mapping is available for the specified port. This is required by ONTAP® software to monitor the health of the hypervisor uplinks.

### Corrective Action

Verify the Hypervisor uplink mapping configuration. Without it, ONTAP cannot monitor the health of the hypervisor uplinks, preventing data LIFs from failing over if the uplink were to go down.

### Syslog Message

The %s health monitor cannot access the Hypervisor uplink mapping configuration for node %s port %s.

### Parameters

**monitor** (STRING): Health monitor name.

**node** (STRING): Name of the node.

**port** (STRING): Name of the port.

## vifmgr.hv.nouplinkstatus

### Severity

ERROR

### Description

This message occurs when there is an error obtaining the Hypervisor uplink status for the specified port.

### Corrective Action

Verify the Hypervisor's host and SNMP configuration.

### Syslog Message

Unable to obtain the hypervisor uplink status for port %s with hypervisor uplinks %s.

### Parameters

**port** (STRING): Name of the port.

**uplinks** (STRING): List of hypervisor uplinks associated with the port.

## vifmgr.ifgrp events

### vifmgr.ifgrp.invalidname

#### Severity

ERROR

#### Description

This message occurs when an interface group (ifgrp) is configured with an unsupported name. This ifgrp cannot be modified or destroyed using clustershell or Data ONTAP® management APIs.

#### Corrective Action

Special corrective action is necessary to remove this ifgrp; contact NetApp technical support for assistance.

#### Syslog Message

Ifgrp %s has an invalid name and cannot be modified or destroyed.

#### Parameters

**ifgrp** (STRING): Invalid name of the ifgrp.

## vifmgr.ipspace events

### vifmgr.ipspace.tooMany

#### Severity

EMERGENCY

#### Description

This message occurs when the system is over-provisioned with IPspaces, exceeding the maximum number allowed. The over-provisioned IPspaces might not be configured successfully in the stack. Any ports assigned to the unconfigured IPspace will fail to host IP addresses and support traffic. This limitation is checked whenever the node comes into quorum.

### Corrective Action

Use the "network ipspace delete" command to remove unused IPspaces.

### Syslog Message

There are %d custom IPspaces configured on the cluster, exceeding the maximum limit of %d.

### Parameters

**configNum** (INT): Number of custom IPspaces configured.

**maxNum** (INT): Maximum number of custom IPspaces supported.

## vifmgr.lif events

### vifmgr.lif.foPolicyMod

#### Severity

NOTICE

#### Description

This message occurs when a LIF's failover policy has been modified.

#### Corrective Action

(None).

#### Syslog Message

LIF %s on Vserver %s (numeric ID: %d, IP address: %s) has been updated to use failover policy %s."

#### Parameters

**lif\_name** (STRING): Name of the affected LIF.

**vserver\_name** (STRING): Vserver associated with the affected LIF.

**lif\_id** (INT): Numeric identifier for the affected LIF.

**lif\_address** (STRING): IP address for the affected LIF.

**new\_policy** (STRING): Failover policy that is now associated with this LIF.

**old\_policy** (STRING): Failover policy that was previously associated with this LIF.

### vifmgr.lif.fwpolicy.update

#### Severity

NOTICE

#### Description

This message occurs when a LIF (logical interface) firewall policy value is deprecated and automatically assigned a new value based on its role.

#### Corrective Action

(None).

#### Syslog Message

LIF %s (on virtual server %s), firewall policy "%s" is deprecated and updated to a new value "%s" based on the LIF's role.



## Parameters

**vifname** (STRING): Name of the LIF whose firewall policy value is updated.

**vsrvid** (STRING): Identifier of the Vserver on which this LIF resides.

**oldpolicy** (STRING): LIF firewall policy value before change.

**newpolicy** (STRING): LIF firewall policy value after change.

## vifmgr.lif.invalidPort

### Severity

ERROR

### Description

This message occurs when a LIF's current or home port are not present in RDB.

### Corrective Action

Use the command in advanced mode "network interface show" making sure to specify the parameter -numeric-id with the identifier specified in the message. With the LIF name and vserver from the output of the previous command, use the "network interface modify" command to set the home port and current port to ports that are displayed by the "network ports show" command.

### Syslog Message

LIF %d has a home port identifier or current port identifier that is not present in RDB.

### Parameters

**lif** (INT): Identifier of LIF with invalid port.

## vifmgr.lif.invalidSubnet

### Severity

ERROR

### Description

This message occurs when a LIF's configuration conflicts with that of its subnet. It is caused by an inconsistency in the configuration of the LIF.

### Corrective Action

To identify the LIF, use the advanced mode command "network interface show" setting the numeric-id parameter to that of the LIF in the message. If the subnet is not set correctly, assign the subnet with the "network interface modify" command specifying an existing subnet that contains the ranges for the LIF's address. If the subnet does not include a range with the address of the LIF, add a range with the "network modify subnet" command making sure to specify the LIF and the force-update-lif-associations parameters.

### Syslog Message

LIF %d has a configuration that does not match its subnet.

### Parameters

**lif** (INT): Identifier of LIF with invalid subnet.

## vifmgr.lif.invFailoverGroup

**Severity**

ERROR

**Description**

This message occurs when a LIF has a failover group with the wrong virtual server or ipspace. If an event occurs where the LIF needs to fail off the current port the LIF might not be hosted.

**Corrective Action**

To identify the LIF, use the advanced mode command "network interface show" setting the numeric-id parameter to that of the LIF in the message. Next, use the "network interface modify" command to associate the LIF with an existing failover group or a new failover group consisting of the desired ports.

**Syslog Message**

The failover group %s for LIF %d is invalid.

**Parameters**

**failovergroup** (STRING): Invalid failover group name.  
**lif** (INT): Identifier of the LIF with an invalid failover group.

**vifmgr.lif.invPortIPs****Severity**

ERROR

**Description**

This message occurs when a port and a LIF currently on that port are in different IPspaces. It is caused by an inconsistency in the configuration of the LIF.

**Corrective Action**

Use the command in advanced mode "network interface show" making sure to specify the paramter -numeric-id with the identifier specified in the message. If the ipspace assigned to the port is correct, modify the LIF to use a home port in the same ipspace as its containing virtual server. If the ipspace of the port is incorrect, change it to match the LIF. If both are correct the LIF's home port, current port, and failover group have to be reassigned to ones sharing the LIF's ipspace.

**Syslog Message**

LIF %d and Port %s:%s are in different IPspaces.

**Parameters**

**lif** (INT): Identifier of LIF with mismatched IpSPACE.  
**node** (STRING): Name of node that port resides on.  
**port** (STRING): Name of port with mismatched IpSPACE.

**vifmgr.lif.orphanedFwPolicy****Severity**

ERROR

**Description**

This message occurs when the virtual server for a firewall policy does not exist.

### Corrective Action

Use the command in advanced mode "network interface show" making sure to specify the parameter -numeric-id with the identifier specified in the message. Then use the command "network interface modify" to assign a new firewall policy with the information from the show command.

### Syslog Message

The firewall policy %s assigned to LIF %d does not exist or has a different virtual server identifier than the LIF.

### Parameters

**policy** (STRING): Firewall policy identifier which has an invalid virtual server.  
**lif** (INT): Firewall policy identifier which has an invalid virtual server.

## vifmgr.lif.protocolMod

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) is temporarily brought offline to add or remove data protocols. Data service will be briefly interrupted while this change is applied, and any existing network connection will be reset.

### Corrective Action

(None).

### Syslog Message

LIF %s on Vserver %s (numeric ID: %d, IP address: %s) is being temporarily brought offline to add or remove data protocols.

### Parameters

**lif\_name** (STRING): Name of the affected LIF.  
**vserver\_name** (STRING): Vserver associated with the affected LIF.  
**lif\_id** (INT): Numeric identifier for the affected LIF.  
**lif\_address** (STRING): IP address for the affected LIF.

## vifmgr.lif.sp.conflict

### Severity

ALERT

### Description

This message occurs when a logical interface (LIF) in the default IPspace is configured with an address that conflicts with the Service Processor (SP) or the Baseboard Management Controller (BMC).

### Corrective Action

Use the "network interface modify" command to assign a different IP address to the affected logical network interface (LIF). Be sure to assign a new address that does not conflict with other hosts in the default IPspace.

## Syslog Message

LIF "%s" on Vserver "%s" with IP address "%s" in the default IPspace is in use by the SP/BMC.

## Parameters

**vifName** (STRING): Name of the LIF that is being moved.  
**vsrvName** (STRING): ID of the virtual server on which the LIF resides.  
**vifip** (STRING): IP address of the affected LIF.

## vifmgr.lif.subnetMisconfig

### Severity

ERROR

### Description

This message occurs when LIFs in the same subnet are configured on ports in multiple broadcast domains. This can cause routing issues due to reachability to and beyond the gateway.

### Corrective Action

To identify the misconfigured LIFs use the command "network interface show -address -netmask -fields broadcast-domain". LIFs in the same subnet should be configured on ports that are in the same broadcast domain.

## Syslog Message

LIFs in subnet %s of IPspace "%s" are configured on ports in multiple broadcast domains: %s

## Parameters

**subnetname** (STRING): Subnet name.  
**ipspace** (STRING): IPspace name.  
**bdname** (STRING): List of broadcast domains.

## vifmgr.lif.svcPolicyMod

### Severity

NOTICE

### Description

This message occurs when a LIF's service policy has been modified.

### Corrective Action

(None).

## Syslog Message

LIF %s on Vserver %s (numeric ID: %d, IP address: %s) has been updated to use service policy %s."

## Parameters

**lif\_name** (STRING): Name of the affected LIF.  
**vserver\_name** (STRING): Vserver associated with the affected LIF.  
**lif\_id** (INT): Numeric identifier for the affected LIF.  
**lif\_address** (STRING): IP address for the affected LIF.  
**new\_policy** (STRING): Service policy that is now associated with this LIF.  
**old\_policy** (STRING): Service policy that was previously associated with this LIF.

## vifmgr.lifadded events

### vifmgr.lifadded.byadmin

#### Severity

NOTICE

#### Description

This message occurs when an administrator action causes a logical interface (LIF) IP address to be added.

#### Corrective Action

(None).

#### Syslog Message

LIF %s with ID %s (on virtual server %s), IP address %s, was added.

#### Parameters

**vifname** (STRING): Name of the LIF that was added.

**vifid** (STRING): Unique ID assigned to the LIF.

**vsrvid** (STRING): Name of the virtual server identifier on which the LIF resides.

**vifip** (STRING): LIF IP address that was added.

## vifmgr.lifbeinghosted events

### vifmgr.lifbeinghosted

#### Severity

NOTICE

#### Description

This message occurs when a logical interface (LIF) IP address that was previously marked as down because no eligible ports could host this LIF is now being hosted. Data access through this LIF is again up.

#### Corrective Action

(None).

#### Syslog Message

LIF %s (on virtual server %s), IP address %s, previously marked as down, is now up and serving data on node %s, port %s.

#### Parameters

**vifname** (STRING): Name of the LIF that is now being hosted.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is now being hosted.

**tonode** (STRING): Name of the node where the LIF is currently being hosted.

**toport** (STRING): Name of the port where the LIF is currently being hosted.

## vifmgr.lifbeingremoved events

## vifmgr.lifBeingRemoved

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) IP address is being removed or moved from the node. Data access through this LIF will be temporarily down until the LIF is re-hosted on another node.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is being removed from node %s, port %s.

### Parameters

**vifname** (STRING): Name of the LIF that is being removed.  
**vsrvid** (STRING): Virtual server identifier on which this LIF resides.  
**vifip** (STRING): The IP address that is being removed.  
**tonode** (STRING): Name of the node on which the LIF had been hosted.  
**toport** (STRING): Name of the port on which the LIF had been hosted.

## vifmgr.lifcapacity events

### vifmgr.lifCapacity

#### Severity

NOTICE

#### Description

This message occurs when the system informs the administrator of the number of Logical Interfaces (LIFs) supported by the node.

#### Corrective Action

(None).

#### Syslog Message

Node %s is capable of hosting %s data LIFs.

#### Parameters

**nodeName** (STRING): Name of the node.  
**numLifs** (STRING): Number of data LIFs that can be configured on the node.

## vifmgr.lifdeleted events

### vifmgr.lifdeleted.byadmin

#### Severity

NOTICE

## Description

This message occurs when an administrator action causes a logical interface (LIF) IP address to be deleted.

## Corrective Action

(None).

## Syslog Message

LIF %s with ID %s (on virtual server %s), IP address %s, was deleted.

## Parameters

**vifname** (STRING): Name of the LIF that was deleted.

**vifid** (STRING): Unique ID assigned to the LIF.

**vsrvid** (STRING): Name of the virtual server identifier on which the LIF resided.

**vifip** (STRING): LIF IP address that was deleted.

# vifmgr.lifdown events

## vifmgr.lifdown.noports

### Severity

ALERT

### Description

This message occurs when a logical interface (LIF) IP address cannot serve data because there are no eligible ports on which to host the LIF. Thus, this LIF is being marked as down. This means that data access through this LIF will be down until the system can rehost this LIF on an eligible port that becomes available.

### Corrective Action

This LIF IP address is down because there are no eligible ports capable of serving data for this LIF. To bring this LIF up so that it can serve data, enable one or more eligible ports that can host this LIF. To find the list of eligible ports for this LIF, examine the failover rules for this LIF. In addition to the ports found there, you could add a new port to the failover rule for that LIF.

### Syslog Message

LIF %s (on virtual server %s), IP address %s, currently cannot be hosted on node %s, port %s, or any of its failover targets, and is being marked as down.

### Parameters

**vifname** (STRING): Name of the LIF that is marked as down.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is marked as down.

**tonode** (STRING): Name of the node to which the LIF currently cannot be hosted.

**toport** (STRING): Name of the port to which the LIF currently cannot be hosted.

# vifmgr.lifidoutofsync events

## vifmgr.lifldOutOfSync

## Severity

ERROR

## Description

This message occurs when a Logical Interface (LIF) is configured in the Logical Interface manager (VIFMgr) but is not visible in the user interface."

## Corrective Action

Using information from the message, use the "network interface ids create" diag privilege command to create the ID.

## Syslog Message

LIF %s with ID %s (on virtual server %s), IP address %s, is configured in the VIFMgr but is not visible in the user interface.

## Parameters

**vifname** (STRING): Name of the invisible LIF as it exists in the VIFMgr.

**vifid** (STRING): Unique ID of the invisible LIF as it exists in VIFMgr.

**vsrvid** (STRING): ID of the virtual server on which the LIF resides.

**vifip** (STRING): IP address of the invisible LIF as it exists in the VIFMgr.

# vifmgr.lifmove events

## vifmgr.lifmove.blocked

## Severity

NOTICE

## Description

This message occurs when an attempt to move a logical interface (LIF) IP address to another node/port due to a load-balance event has failed because the LIF has active connections which can not be safely migrated.

## Corrective Action

(None).

## Syslog Message

LIF %s (on virtual server %s), IP address %s, was blocked when being moved to node %s, port %s.

## Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

# vifmgr.lifmoved events



## vifmgr.lifmoved.byadmin

### Severity

NOTICE

### Description

This message occurs when an administrator action causes a logical interface (LIF) IP address to move to another node/port.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is being moved to node %s, port %s.

### Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

## vifmgr.lifmoved.bylb

### Severity

NOTICE

### Description

This message occurs when an attempt is made to move a logical interface (LIF) IP address to another node/port due to a load-balance event.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is being moved to node %s, port %s due to an LB event.

### Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

## vifmgr.lifmoved.linkdown

### Severity

NOTICE

### Description

This message occurs when an attempt was made to move a logical interface (LIF) IP address to another node/port due to a link down event.

## Corrective Action

(None).

## Syslog Message

LIF %s (on virtual server %s), IP address %s, is being moved to node %s, port %s.

## Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

## vifmgr.lifmoved.nodedown

### Severity

NOTICE

### Description

This message occurs when an attempt was made to move a logical interface (LIF) IP address to another node/port due to a node down event.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is being moved to node %s, port %s.

### Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

## vifmgr.lifondegport events

### vifmgr.lifOnDegPort.noPorts

### Severity

ALERT

### Description

This message occurs when an attempt to move a logical interface (LIF) currently hosted on a port with degraded health status to a healthy port fails because there are no healthy ports in the LIF's failover group.

### Corrective Action

Add ports to the broadcast domain or failover group used by the affected LIF. Use the "network interface show -failover" command to review the currently assigned failover targets for each LIF.

## Syslog Message

LIF %s (on virtual server %s), IP address %s, is currently hosted on node %s, port %s, whose health status is degraded.

## Parameters

**vifName** (STRING): Name of the LIF that is hosted on the port with degraded health status.  
**vsrvId** (STRING): ID of the virtual server on which the LIF resides.  
**vifIp** (STRING): IP address of the LIF that is hosted on the port with degraded health status.  
**node** (STRING): Name of the node on which the degraded port resides.  
**port** (STRING): Name of the port with degraded health status.

# vifmgr.lifs events

## vifmgr.lifs.clusterlimit

### Severity

ERROR

### Description

This message occurs when the total number of logical interfaces (LIFs) exceeds the supported per-cluster limit. To ensure system stability, the total number of IP data LIFs should be kept at or below this limit.

### Corrective Action

Use the "network interface show -role data" command to review existing LIFs, and use the "network interface delete" command to remove enough LIFs to meet the cluster maximum. To facilitate this process, you can display the current number of LIFs by using the "network interface capacity show" command.

## Syslog Message

The total number of data LIFs in this cluster (%u) exceeds the supported limit of %u.

## Parameters

**numLifs** (INT): The number of IP data LIFs present in this cluster.  
**lifLimit** (INT): The maximum number of supported LIFs for this cluster.

## vifmgr.lifs.lowredundancy

### Severity

ALERT

### Description

This message occurs when one or more logical interfaces (LIFs) are configured to use a failover policy that implies failover to a specific set of nodes but does not have available failover targets on enough of those nodes. If the specified home node is offline or unavailable, all affected LIFs will be operationally down and unable to serve any data.

### Corrective Action

Add additional ports to the broadcast domains or failover groups used by the affected LIFs, or modify each LIF's failover policy to include one or more nodes with available failover targets. For example, the "broadcast-domain-wide" failover policy will consider all failover targets in a LIF's failover group. Use the "network interface show -failover" command to review the currently assigned failover targets for each LIF.

## Syslog Message

Insufficient redundancy in the failover configuration for %u LIFs assigned to node "%s". LIFs: %s

### Parameters

**lifcount** (INT): Number of affected LIFs.  
**homenode** (STRING): Home node of the affected LIFs.  
**liflist** (STRING): List of the first several affected LIFs.

## vifmgr.lifs.noredundancy

### Severity

ALERT

### Description

This message occurs when one or more logical interfaces (LIFs) are configured to use a failover policy that implies failover to one or more ports but have no failover targets beyond their home ports. If any affected home port or home node is offline or unavailable, the corresponding LIFs will be operationally down and unable to serve data.

### Corrective Action

Add additional ports to the broadcast domains or failover groups used by the affected LIFs, or modify each LIF's failover policy to include one or more nodes with available failover targets. For example, the "broadcast-domain-wide" failover policy will consider all failover targets in a LIF's failover group. Use the "network interface show -failover" command to review the currently assigned failover targets for each LIF. If the intent is to never fail over the LIF to any other failover targets, set the failover-policy for the LIF to 'disabled' to prevent this alert in the future. For additional information about configuring LIF failover targets, search the NetApp® support site for TR-4182.

## Syslog Message

No redundancy in the failover configuration for %u LIFs assigned to node "%s". LIFs: %s

### Parameters

**lifcount** (INT): Number of affected LIFs.  
**homenode** (STRING): Home node of the affected LIFs.  
**liflist** (STRING): List of the first several affected LIFs.

## vifmgr.lifs.overcapacity

### Severity

ALERT

### Description

This message occurs when logical interfaces (LIFs) have been over-provisioned. As a result, some LIFs might report a status of "down" during a node outage, even if the LIF is configured to have a remote failover target.

### Corrective Action

Removing LIFs from the system, adjusting the failover settings of some LIFs, and reconfiguring failover groups are all possible options for correcting the situation. Use the "network interface check failover show" command to gather more detailed information about the over-provisioning.

## Syslog Message

Node with failover targets over-provisioned: %u LIFs could be marked "down" during node outage.

## Parameters

**overamount** (INT): Number of LIFs that might be displayed as "down".

# vifmgr.lifsuccessfullymoved events

## vifmgr.lifsuccessfullymoved

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) IP address was successfully moved to another node/port.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is now hosted on node %s, port %s.

### Parameters

**vifname** (STRING): Name of the LIF that was moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that was moved.

**tonode** (STRING): Name of the node to which the LIF was moved.

**toport** (STRING): Name of the port to which the LIF was moved.

# vifmgr.migratelifs events

## vifmgr.migrateLifs.nlbdDown

### Severity

ERROR

### Description

This message occurs when the network load balancing listener is not responding to probes over a node management LIF, and cluster-scoped LIFs with the same current port have been migrated elsewhere.

### Corrective Action

Use the "network interface show" command to verify that the node management LIF in question is operational. If the LIF is operational, contact NetApp technical support for specific instructions.

### Syslog Message

The network load balancing listener is not responding via LIF "%s" on Vserver "%s" with current port "%s". As a result, the load balancing LIFs have been migrated elsewhere.

## Parameters

**lif** (STRING): LIF name.  
**vserver\_name** (STRING): Vserver name.  
**port** (STRING): Network port name

## vifmgr.migrateLifs.nlbdUp

### Severity

NOTICE

### Description

This message occurs when the network load balancing listener starts responding to probes over a node management LIF. The network port where the node management LIF is located is now available for use by load-balanced LIFs.

### Corrective Action

(None).

### Syslog Message

The network load balancing listener has started responding via LIF "%s" on Vserver "%s" with current port "%s".

## Parameters

**lif** (STRING): LIF name.  
**vserver\_name** (STRING): Vserver name.  
**port** (STRING): Network port name.

## vifmgr.port events

### vifmgr.port.discovered

#### Severity

NOTICE

#### Description

This message occurs when a physical network port has been discovered.

#### Corrective Action

(None).

#### Syslog Message

Physical network port %s (node %s) has been discovered.

#### Parameters

**port** (STRING): Name of the network port.  
**node** (STRING): Name of the node on which the network port resides.

### vifmgr.port.ifgrpAddPort

**Severity**

NOTICE

**Description**

This message occurs when an administrative action adds a port to an interface group.

**Corrective Action**

(None).

**Syslog Message**

Port %s, node %s has been added to ifgrp %s.

**Parameters**

**portName** (STRING): Name of the port that was added to an interface group.

**nodeName** (STRING): Name of the node on which the port resides.

**ifgrpName** (STRING): Name of the interface group to which the port was added.

**vifmgr.port.ifgrpCreated****Severity**

NOTICE

**Description**

This message occurs when an administrative action creates an interface group.

**Corrective Action**

(None).

**Syslog Message**

Ifgrp %s, node %s has been created.

**Parameters**

**ifgrpName** (STRING): Name of the interface group that was created.

**nodeName** (STRING): Name of the node on which the interface group resides.

**vifmgr.port.ifgrpDeleted****Severity**

NOTICE

**Description**

This message occurs when an administrative action deletes an interface group.

**Corrective Action**

(None).

**Syslog Message**

Ifgrp %s, on node %s has been deleted.

## Parameters

**ifgrpName** (STRING): Name of the interface group that was deleted.  
**nodeName** (STRING): Name of the node on which the ifgrp was deleted.

## vifmgr.port.ifgrpRemovePort

### Severity

NOTICE

### Description

This message occurs when an administrative action removes a port from an interface group.

### Corrective Action

(None).

### Syslog Message

Port %s, node %s has been removed from ifgrp %s.

## Parameters

**portName** (STRING): Name of the port that was removed from an interface group.  
**nodeName** (STRING): Name of the node on which the port resides.  
**ifgrpName** (STRING): Name of the interface group from which the port was removed.

## vifmgr.port.mismatchState

### Severity

ERROR

### Description

This message occurs when an administratively down port is not operationally down.

### Corrective Action

The operational state of the port does not match the state in the ports configuration. Use the advanced mode command "network port modify" command specifying the desired state by setting the paramter -up -admin" to true if the desired state is up and to false if the desired state is down.

### Syslog Message

port %s is administratively down and operationally up.

## Parameters

**port** (STRING): Name of port that should be operationally down.

## vifmgr.port.monitor.failed

### Severity

ERROR

### Description

This message occurs when a port fails a health check.



## Corrective Action

Check the network cables that are connected to the specified port for any loose connections. Ensure that the physical Ethernet NIC is functioning properly.

## Syslog Message

The "%s" health check for port %s (node %s) has failed. The port is operating in a degraded state.

## Parameters

**healthCheckType** (STRING): The type of health check that this port failed. The possible health checks are Frequent Link Flapping and Insufficient L2 Reachability.

**portName** (STRING): Name of the port that failed the health check.

**nodeName** (STRING): Name of the node on which the port resides.

## vifmgr.port.monitor.passed

### Severity

NOTICE

### Description

This message occurs when a port successfully passes a health check.

### Corrective Action

(None).

### Syslog Message

The "%s" health check for port %s (node %s) has passed.

### Parameters

**healthCheckType** (STRING): The type of health check that this port passed. The possible health checks are Frequent Link Flapping and Insufficient L2 Reachability.

**portName** (STRING): Name of the port that passed the health check.

**nodeName** (STRING): Name of the node on which the port resides.

## vifmgr.port.mtumismatch

### Severity

ERROR

### Description

This message occurs when the administrator modifies a cluster port's MTU while the node is out of quorum to a value other than the broadcast domain MTU, and then the node rejoins quorum.

### Corrective Action

Change the port's MTU by using the "network port modify" command, or change the broadcast domain's MTU by using the "network port broadcast-domain modify" command.

### Syslog Message

The MTU of port %s (%d) does not match the MTU of broadcast domain %s (%d).

### Parameters

**port** (STRING): Name of the port.

**portmtu** (INT): MTU of the port.

**broadcast\_domain** (STRING): Name of the broadcast domain.

**bdmtu** (INT): MTU of the broadcast domain.

## vifmgr.port.restored

### Severity

NOTICE

### Description

This message occurs when configuration data for a network port is restored, typically for a network port that no longer exists or was partially removed previously. No action is required unless you were intentionally intending to remove the network port.

### Corrective Action

If you intentionally intended to remove the port, use the (privilege: advanced) "cluster ring show -unitname vifmgr" command to confirm that the network management component is online. If the component is online, use either the "network port ifgrp delete" or "network port vlan delete" command to remove the port again.

### Syslog Message

Configuration data for port %s in IPspace "%s" has been restored.

### Parameters

**port** (STRING): Name of the affected port.

**ipspace** (STRING): Name of the IPspace to which the port is assigned.

## vifmgr.port.shareBd

### Severity

ERROR

### Description

This message occurs when an administrator attempts to move a network port out of IPspace "Cluster", and that port is being used by the back-end HA or DR interconnect. Ports used for either the HA or DR interconnect must remain in IPspace "Cluster" and must be wired to either the cluster switch or directly to the associated HA partner, for two-node switchless clusters.

### Corrective Action

If the port has been added to a new broadcast domain, use the "network port broadcast-domain remove-ports" command to remove the port's broadcast domain association. When the port is no longer assigned to a broadcast domain, use the "network port broadcast-domain add-ports" command to add it to broadcast domain "Cluster" in IPspace "Cluster".

### Syslog Message

Port %s is shared with the HA or DR interconnect but has been moved from broadcast domain "Cluster" to broadcast domain "%s" of IPspace "%s".

### Parameters

**port** (STRING): Name of the affected port.

**broadcast\_domain** (STRING): Name of the broadcast domain to which the port is assigned.

**ipspace** (STRING): Name of the IPspace to which the port is assigned.

## vifmgr.port.sharelpspace

### Severity

ERROR

### Description

This message occurs when a cluster port that is shared with either the back-end HA or DR interconnect is assigned to an IPspace other than "Cluster", which is required for all shared ports. This condition can be encountered when a hardware refresh introduces a new platform that includes ports shared with either the HA or DR interconnect.

### Corrective Action

If necessary, use the "network port broadcast-domain remove-ports" command to remove the broadcast domain association from the port. After the port is no longer assigned to a broadcast domain, use the "network port broadcast-domain add-ports" command to add the port to broadcast domain "Cluster" in IPspace "Cluster".

### Syslog Message

Port %s in broadcast domain "%s" of IPspace "%s" is shared with the HA or DR interconnect, but is not assigned to IPspace "Cluster".

### Parameters

**port** (STRING): Name of the affected port.

**broadcast\_domain** (STRING): Name of the broadcast domain to which the port is assigned.

**ipspace** (STRING): Name of the IPspace to which the port is assigned.

## vifmgr.port.shareMtu

### Severity

ERROR

### Description

This message occurs when a cluster port shared with either the back-end HA or DR interconnect is found to be using an MTU other than 9000 bytes. For optimal performance, ports shared with the back-end interconnect should use an MTU of 9000 bytes. This condition can occur when a platform using network ports shared with either the back-end HA or DR interconnect is used in a mixed-version cluster, where one or more nodes still allows changes to the MTU setting for broadcast domain "Cluster".

### Corrective Action

If the port is currently assigned to broadcast domain "Cluster", use the "network port broadcast-domain modify" command to update broadcast domain "Cluster" to use an MTU of 9000 bytes.

### Syslog Message

Port %s in broadcast domain "%s" of IPspace "%s" is shared with the HA or DR interconnect, but has an MTU of %d bytes instead of 9000 bytes.

### Parameters

**port** (STRING): Name of the affected port.

**broadcast\_domain** (STRING): Name of the broadcast domain to which the port is assigned.

**ipspace** (STRING): Name of the IPspace to which the port is assigned.

**mtu** (INT): MTU setting of the port.

## **vifmgr.port.sharingDisabled**

### **Severity**

NOTICE

### **Description**

This message occurs when port sharing is disabled for a network port. This transition occurs when a network port is no longer available to be shared with the back-end HA or DR interconnect.

### **Corrective Action**

(None).

### **Syslog Message**

Sharing is disabled for port %s.

### **Parameters**

**port** (STRING): Name of the affected port.

## **vifmgr.port.sharingEnabled**

### **Severity**

NOTICE

### **Description**

This message occurs when port sharing is enabled for a network port. Enabling port sharing facilitates bandwidth sharing with the back-end (high availability) HA and DR interconnect, which use the same physical port.

### **Corrective Action**

(None).

### **Syslog Message**

Sharing is enabled for port %s.

### **Parameters**

**port** (STRING): Name of the affected port.

## **vifmgr.port.sharingError**

### **Severity**

ERROR

### **Description**

This message occurs when an error is encountered while attempting to enable or disable sharing of a network port. If the port is being shared with the back-end HA or DR interconnect on this platform, those components might not be safely isolated from LIFs configured on the port. High utilization by either HA or DR traffic might negatively impact latency and bandwidth available to the cluster network.

### **Corrective Action**

Use the "network port show" command to display the current link status of the affected network port. If the port is down, check any associated cabling and the status of the network switch it is connected to. If the problem persists, contact NetApp technical support for assistance.

## Syslog Message

Port "%s" could not be updated to the requested sharing state. Error: %s.

## Parameters

**port** (STRING): Name of the affected port.

**error** (STRING): Description of the error.

## vifmgr.port.vipCreated

### Severity

NOTICE

### Description

This message occurs when an administrative action causes a VIP port to be created.

### Corrective Action

(None).

## Syslog Message

VIP port %s on node %s in IPspace %s has been created.

## Parameters

**vipPortName** (STRING): Name of the VIP port that was created.

**nodeName** (STRING): Name of the node on which the VIP port resides.

**ipspace** (STRING): Name of the IPspace where the VIP port resides.

## vifmgr.port.vipDeleted

### Severity

NOTICE

### Description

This message occurs when an administrative action causes a VIP port to be deleted.

### Corrective Action

(None).

## Syslog Message

VIP port %s on node %s in IPspace %s has been deleted.

## Parameters

**vipPortName** (STRING): Name of the VIP port that was deleted.

**nodeName** (STRING): Name of the node from which the VIP port was deleted.

**ipspace** (STRING): Name of the IPspace where the VIP port resides.

## vifmgr.port.vlanCreated

### Severity

NOTICE

## Description

This message occurs when an administrative action creates a VLAN on a port.

## Corrective Action

(None).

## Syslog Message

Vlan %s is created on port %s, node %s.

## Parameters

**vlanName** (STRING): Name of the VLAN created on the port.

**portName** (STRING): Name of the port on which the VLAN was created.

**nodeName** (STRING): Name of the node on which the port resides.

## vifmgr.port.vlanDeleted

### Severity

NOTICE

### Description

This message occurs when an administrative action deletes a VLAN from a port.

### Corrective Action

(None).

### Syslog Message

Vlan %s is deleted from port %s, node %s.

### Parameters

**vlanName** (STRING): Name of the VLAN deleted from the port.

**portName** (STRING): Name of the port from which the VLAN was deleted.

**nodeName** (STRING): Name of the node on which the port resides.

## vifmgr.portdeg events

### vifmgr.portDeg.ifgrpAdded

#### Severity

NOTICE

#### Description

This message occurs when an administrative action adds a port with a degraded health status to an interface group.

#### Corrective Action

(None).

#### Syslog Message

Port %s, node %s whose health-status is degraded has been added to ifgrp %s.

## Parameters

- portName** (STRING): Name of the degraded port that was added to an interface group.
- nodeName** (STRING): Name of the node on which the degraded port resides.
- ifgrpName** (STRING): Name of the interface group to which the degraded port was added.

## vifmgr.portDeg.lifHosted

### Severity

NOTICE

### Description

This message occurs when an administrative action hosts a logical interface (LIF) on a port with a degraded health status.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s) is hosted on a degraded port %s, node %s.

## Parameters

- vifName** (STRING): Name of the hosted LIF.
- vsrvid** (STRING): Name of the virtual server identifier that uses the LIF.
- portName** (STRING): Name of the degraded port that hosts the LIF.
- nodeName** (STRING): Name of the node on which the degraded port resides.

## vifmgr.portDeg.lifMoved

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) is moved from a port with degraded health status to a healthy port.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, is being moved to node %s, port %s from node %s, port %s because the port's health status is in a degraded state.

## Parameters

- vifName** (STRING): Name of the LIF that is being moved.
- vsrvid** (STRING): ID of the virtual server on which the LIF resides.
- vifIp** (STRING): The IP address of the LIF that is being moved.
- toNode** (STRING): Name of the node to which the LIF is moving.
- toPort** (STRING): Name of the port to which the LIF is moving.
- fromNode** (STRING): Name of the node on which the LIF is currently hosted.
- fromPort** (STRING): Name of the port on which the LIF is currently hosted.

## vifmgr.portDeg.vlanCreated

### Severity

NOTICE

### Description

This message occurs when an administrative action creates a VLAN on a port with a degraded health status.

### Corrective Action

(None).

### Syslog Message

Vlan %s is created on a degraded port %s, node %s.

### Parameters

**vlanName** (STRING): Name of the VLAN created on the degraded port.

**portName** (STRING): Name of the degraded port on which the VLAN was created.

**nodeName** (STRING): Name of the node on which the degraded port resides.

## vifmgr.portdown events

### vifmgr.portdown

#### Severity

NOTICE

#### Description

This message occurs when a link down event was received on a network port. Data access through this physical port is down.

#### Corrective Action

(None).

#### Syslog Message

A link down event was received on node %s, port %s.

#### Parameters

**nodename** (STRING): Name of the node on which the affected port resides.

**portname** (STRING): Name of the affected port.

## vifmgr.porthhealthy events

### vifmgr.portHealthy

#### Severity

NOTICE



## Description

This message occurs when port health status changes from degraded to healthy.

## Corrective Action

(None).

## Syslog Message

The health status of port %s (node %s) has changed from degraded to healthy. This port can now be used for hosting LIFs.

## Parameters

**portName** (STRING): Name of the port whose health status changed from degraded to healthy.

**nodeName** (STRING): Name of the node on which the port resides.

# vifmgr.portowner events

## vifmgr.portowner.err

### Severity

ERROR

### Description

This message occurs when a port receives a configuration error.

### Corrective Action

Remove any ports from the bootarg.bsdportname LOADER variable that are of a role other than node-mgmt by resetting the bootarg.bsdportname under the LOADER prompt: LOADER> setenv bootarg.bsdportname "port-name (node-mgmt port only)" Then reboot the node.

### Syslog Message

Node %s port %s has an invalid port role type. Only node-mgmt network ports can be assigned to use a different driver.

### Parameters

**nodename** (STRING): Name of the node on which the affected port resides.

**portname** (STRING): Name of the affected port.

# vifmgr.portremoved events

## vifmgr.portRemoved

### Severity

NOTICE

### Description

This message occurs when a port is removed because it no longer exists.

### Corrective Action

(None).

## Syslog Message

Port %s on node %s has been removed because it no longer exists.

## Parameters

**port** (STRING): Name of the nonexistent port.

**node** (STRING): Name of the node on which the nonexistent port used to reside.

## vifmgr.portRemoved.lifHomed

### Severity

NOTICE

### Description

This message occurs when a logical interface's (LIF's) home port is updated because its previous home port has been removed from its broadcast domain.

### Corrective Action

(None).

## Syslog Message

Home port of LIF %s (Vserver %s), IP addr %s updated to node %s port %s from node %s port %s: port removed from broadcast domain.

## Parameters

**vifName** (STRING): Name of the LIF on which the home port is being updated.

**vsrvId** (STRING): ID of the virtual server on which the LIF resides.

**vifIp** (STRING): IP address of the LIF on which the home port is being updated.

**toNode** (STRING): Name of the node to which the LIF's home port is being updated.

**toPort** (STRING): Name of the port to which the LIF's home port is being updated.

**fromNode** (STRING): Name of the node from which the LIF's home port is being updated.

**fromPort** (STRING): Name of the port from which the LIF's home port is being updated.

## vifmgr.portRemoved.lifMoved

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) is moved from a port that has been removed from its broadcast domain.

### Corrective Action

(None).

## Syslog Message

LIF %s (Vserver %s), IP addr %s moving to node %s port %s from node %s port %s: port removed from broadcast domain.

## Parameters

**vifName** (STRING): Name of the LIF that is being moved.

**vsrvId** (STRING): ID of the virtual server on which the LIF resides.

**viflp** (STRING): IP address of the LIF that is being moved.  
**toNode** (STRING): Name of the node to which the LIF is moving.  
**toPort** (STRING): Name of the port to which the LIF is moving.  
**fromNode** (STRING): Name of the node from which the LIF is moving.  
**fromPort** (STRING): Name of the port from which the LIF is moving.

## vifmgr.portspeed events

### vifmgr.portspeed.err

#### Severity

ERROR

#### Description

This message occurs when a port is connected to a switch utilizing a link speed that is not qualified. This might lead to unexpected system behaviors.

#### Corrective Action

Ensure that the affected network port is connected to a switch that uses a link speed supported by this node. If necessary, adjust the switch configuration to assign a lower speed.

#### Syslog Message

Node %s port %s is reporting a speed %s that is not supported.

#### Parameters

**nodename** (STRING): Name of the node on which the affected port resides.  
**portname** (STRING): Name of the affected port.  
**portspeed** (STRING): Operational speed of the affected port.

## vifmgr.portup events

### vifmgr.portup

#### Severity

NOTICE

#### Description

This message occurs when a link up event was received on a network port. Data access through this physical port is now available.

#### Corrective Action

(None).

#### Syslog Message

A link up event was received on node %s, port %s.

#### Parameters

**nodename** (STRING): Name of the node on which the affected port resides.  
**portname** (STRING): Name of the affected port.

# vifmgr.rdma events

## vifmgr.RDMA.lifMoved

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) is moved from a port that does not support the required remote direct memory access (RDMA) protocols of the LIF.

### Corrective Action

(None).

### Syslog Message

LIF "%s" (on virtual server "%s"), with IP address "%s", is being moved to node "%s", port "%s", from node "%s", port "%s", because the port does not support the LIF's RDMA protocols.

### Parameters

**vifName** (STRING): Name of the LIF that is being moved.  
**vsrvId** (STRING): ID of the virtual server on which the LIF resides.  
**vifIp** (STRING): IP address of the LIF that is being moved.  
**toNode** (STRING): Name of the node to which the LIF is being moved.  
**toPort** (STRING): Name of the port to which the LIF is being moved.  
**fromNode** (STRING): Name of the node on which the LIF is currently hosted.  
**fromPort** (STRING): Name of the port on which the LIF is currently hosted.

# vifmgr.reach events

## vifmgr.reach.avail

### Severity

NOTICE

### Description

This message occurs when the node has completed its initial scan of reachability for all ports.

### Corrective Action

(None).

### Syslog Message

Node %s has completed its initial scan of reachability for all ports.

### Parameters

**node** (STRING): Name of the node on which the network port resides.

## vifmgr.reach.err

### Severity

ERROR

## Description

This message occurs when a network port encountered errors while scanning for reachability.

## Corrective Action

Verify the network port's physical connectivity and switch configuration. A network port should have reachability only to its expected broadcast domain.

## Syslog Message

Network port %s on node %s encountered errors while scanning for reachability. Reachability was detected to the following broadcast domains: %s.

## Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**reachableBroadcastDomains** (STRING): Names of the broadcast domains that are reachable from the port.

## vifmgr.reach.miscfg

### Severity

ERROR

### Description

This message occurs when a network port cannot reach its expected broadcast domain but has reachability to another broadcast domain.

### Corrective Action

Verify the network port's physical connectivity and switch configuration. A network port should have reachability only to its expected broadcast domain.

### Syslog Message

Network port %s on node %s cannot reach its expected broadcast domain %s, but reachability was detected to broadcast domain %s.

### Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**broadcastDomain** (STRING): Name of the broadcast domain associated with the port.

**reachableBroadcastDomains** (STRING): Names of the broadcast domains that are reachable from the port.

## vifmgr.reach.miscfg.unassn

### Severity

NOTICE

### Description

This message occurs when a network port is not assigned to a broadcast domain but has reachability to a broadcast domain.

### Corrective Action

Verify the network port's physical connectivity and switch configuration. A network port should have reachability only to its expected broadcast domain.

### Syslog Message

Network port %s on node %s is not assigned to a broadcast domain, but reachability was detected to broadcast domain %s.

### Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**reachableBroadcastDomains** (STRING): Names of the broadcast domains that are reachable from the port.

## vifmgr.reach.multireach

### Severity

ERROR

### Description

This message occurs when a network port is determined to have reachability to multiple broadcast domains.

### Corrective Action

Verify the network port's physical connectivity and switch configuration. A network port should have reachability only to its expected broadcast domain.

### Syslog Message

Network port %s on node %s can reach the following broadcast domains: %s.

### Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**reachableBroadcastDomains** (STRING): Names of the broadcast domains that are reachable from the port.

## vifmgr.reach.noreach

### Severity

NOTICE

### Description

This message occurs when a network port cannot reach its expected broadcast domain or any other broadcast domain.

### Corrective Action

Verify the network port's physical connectivity and switch configuration. A network port should have reachability only to its expected broadcast domain.

### Syslog Message

Network port %s on node %s cannot reach its expected broadcast domain %s. No other broadcast domains appear to be reachable from this port.

## Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**broadcastDomain** (STRING): Name of the broadcast domain associated with the port.

## vifmgr.reach.ok

### Severity

NOTICE

### Description

This message occurs when a network port's actual broadcast domain reachability matches its expected reachability, and no other broadcast domains appear to be reachable from the port.

### Corrective Action

(None).

### Syslog Message

Network port %s on node %s can reach its expected broadcast domain %s. No other broadcast domains appear to be reachable from this port.

## Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

**broadcastDomain** (STRING): Name of the broadcast domain associated with the port.

## vifmgr.reach.skipped

### Severity

NOTICE

### Description

This message occurs when a network port was not scanned for reachability because it was administratively or operationally down at the time of the scan.

### Corrective Action

(None).

### Syslog Message

Network port %s on node %s was not scanned for reachability because it was administratively or operationally down at the time of the scan.

## Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

## vifmgr.reach.unassn

### Severity

NOTICE

## Description

This message occurs when a network port is not assigned to a broadcast domain, and has no reachability to any broadcast domains.

## Corrective Action

(None).

## Syslog Message

Network port %s on node %s is not assigned to any broadcast domain. No broadcast domains appear to be reachable from this port.

## Parameters

**port** (STRING): Name of the network port.

**node** (STRING): Name of the node on which the network port resides.

# vifmgr.rpc events

## vifmgr.rpc.nblade.timeouts

### Severity

ERROR

### Description

This message occurs when the virtual interface manager (vifmgr) is having trouble communicating with the packet and connection processing module of the network component. If this problem persists, interface failover might not occur as intended, resulting in a disruption.

### Corrective Action

If the node is experiencing heavy load, the condition might clear up automatically after a few minutes. If the condition persists, contact NetApp technical support.

### Syslog Message

The virtual interface manager is not receiving responses from the node's networking component.

### Parameters

(None).

# vifmgr.secondaryoffline events

## vifmgr.secondaryOffline

### Severity

NOTICE

### Description

This message occurs when the system determines that a remote node has been offline for more than a minute and assigns logical network interfaces (LIFs) from that node to other available nodes in the cluster.



## Corrective Action

(None).

## Syslog Message

LIFs from offline node "%s" are being assigned to other available nodes in the cluster.

## Parameters

**offlineNode** (STRING): Name of the offline node.

# vifmgr.skippingl2pingtest events

## vifmgr.skippingL2PingTest

### Severity

NOTICE

### Description

This message occurs when the number of local ports in the given broadcast domain exceeds the threshold for a node and, as a result, the `I2_reachability` health monitor has omitted checking ports in that broadcast domain. If any local ports in this broadcast domain experience Layer 2 connectivity problems, they are not marked as degraded and no EMS messages related to the Layer 2 connectivity health of those ports is generated. LIFs hosted on those port are not automatically migrated to healthy ports because of Layer 2 connectivity issues.

### Corrective Action

This is not a misconfiguration; however, including fewer than the specified threshold ports in the broadcast domain enables the `I2_reachability` health monitor to check the health of the ports in that broadcast domain.

### Syslog Message

Skipping `I2_reachability` health monitor check for ports belonging to broadcast domain %s on node %s as the number of local ports (%u) belonging to this broadcast domain exceeds the threshold (%u) for which `I2_reachability` health monitor check can efficiently run.

### Parameters

**bcDomainName** (STRING): Name of the broadcast domain for which the `I2_reachability` health monitor check is being skipped.

**nodeName** (STRING): Name of the node for which the number of local ports in the given broadcast domain exceeds the threshold.

**numLocalPorts** (INT): Number of local ports in the specified broadcast domain.

**maxLocalPorts** (INT): Threshold value for the maximum number of local ports in a broadcast domain for which the `I2_reachability` health monitor check can run.

# vifmgr.started events

## vifmgr.started

### Severity

NOTICE

**Description**

This message occurs to notify you that the Logical Interface Manager (VIFMgr) started.

**Corrective Action**

(None).

**Syslog Message**

The Logical Interface Manager (VIFMgr) has started.

**Parameters**

(None).

## vifmgr.startup events

### vifmgr.startup.failover.err

**Severity**

ALERT

**Description**

This message occurs when the logical interface manager (VIFMgr) encounters an error while initializing the structures required for logical interface (LIF) failover. As a result of this error, the node is now running in a degraded state.

**Corrective Action**

Contact NetApp technical support and reference the "vifmgr.startup.failover.err" error.

**Syslog Message**

VIFMgr encountered errors during startup.

**Parameters**

(None).

### vifmgr.startup.merge.err

**Severity**

ERROR

**Description**

This message occurs to notify you that the Logical Interface Manager (VIFMgr) encountered an error while merging the local node's network configuration with the configuration stored in the cluster database. As a result of this error, the node is now running in a degraded state.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

The Logical Interface Manager (VIFMgr) encountered errors during startup.

## Parameters

(None).

# vifmgr.staticroute events

## vifmgr.staticRoute.add

### Severity

NOTICE

### Description

This message occurs when a static route is added to a Vserver.

### Corrective Action

(None).

### Syslog Message

Static route created on Vserver %s: dest %s, gw %s, metric %u.

### Parameters

**vserverName** (STRING): Name of the Vserver on which the static route was added.

**destination** (STRING): The IP subnet to which the route applies.

**gateway** (STRING): The IP address of the gateway.

**metric** (INT): The metric of the route, used to determine its priority relative to other routes for the Vserver.

## vifmgr.staticRoute.del

### Severity

NOTICE

### Description

This message occurs when a static route is deleted from a Vserver.

### Corrective Action

(None).

### Syslog Message

Static route deleted from Vserver %s: dest %s, gw %s, metric %u.

### Parameters

**vserverName** (STRING): Name of the Vserver from which the static route was deleted.

**destination** (STRING): The IP subnet to which the route applies.

**gateway** (STRING): The IP address of the gateway.

**metric** (INT): The metric of the route, used to determine its priority relative to other routes for the Vserver.

# vifmgr.subnet events

## vifmgr.subnet.addr.acquired

**Severity**

NOTICE

**Description**

This message occurs when an address is successfully acquired from a subnet's available-ranges list.

**Corrective Action**

(None).

**Syslog Message**

Address %s has been successfully acquired from subnet %s (ipSPACE %s).

**Parameters**

**address** (STRING): Address acquired.  
**subnet** (STRING): Name of the subnet.  
**ipSPACE** (STRING): Name of the subnet's IPspace.

**vifmgr.subnet.addr.released****Severity**

NOTICE

**Description**

This message occurs when an address is successfully released back to a subnet's available-ranges list.

**Corrective Action**

(None).

**Syslog Message**

Address %s has been successfully released back to subnet %s (ipSPACE %s).

**Parameters**

**address** (STRING): Address released.  
**subnet** (STRING): Name of the subnet.  
**ipSPACE** (STRING): Name of the subnet's IPspace.

**vifmgr.subnet.usedAddr****Severity**

ERROR

**Description**

This message occurs when a LIF's address conflicts with a subnet's list of used and available addresses due to an inconsistency in the configuration of the subnet or LIF.

**Corrective Action**

Either the address of the LIF, the address ranges, or the available addresses in the subnet need to be adjusted. To identify the LIF, use the advanced-privilege command "network interface show" setting the "numeric-id" parameter to that of the LIF in the message. To change the address of the LIF, use the "network interface modify" command, specifying the address and netmask. To remove the address from the subnet, use the "network subnet remove-ranges" command.

## Syslog Message

LIF %d has an address %s that is in Subnet's %s unused pool.

## Parameters

**lif** (INT): Identifier of the LIF with the address conflict.  
**address** (STRING): Address owned by both LIF and subnet.  
**subnet** (STRING): Name of the subnet.

# vifmgr.svcpolicy events

## vifmgr.svcPolicy.addSvc

### Severity

NOTICE

### Description

This message occurs when a service is added to a network interface service policy.

### Corrective Action

(None).

## Syslog Message

Service %s has been added to network interface service policy %s on Vserver %s.

## Parameters

**service** (STRING): Name of the service that has been added.  
**policy** (STRING): Name of the affected service policy.  
**vserver** (STRING): Name of the Vserver associated with this policy.  
**allowed\_addresses** (STRING): List of allowed address masks for this service.

## vifmgr.svcPolicy.deprecated

### Severity

NOTICE

### Description

This message occurs when a network interface (LIF) is configured to use a deprecated service policy.

### Corrective Action

Use the recommended service policy to replace the deprecated service policy.

## Syslog Message

LIF "%s" on Vserver "%s" is using deprecated service policy: %s. Recommended service policy: %s.

## Parameters

**lif\_name** (STRING): LIF name.  
**vserver** (STRING): Vserver name.  
**deprecated\_service\_policy** (STRING): Deprecated service policy name.  
**service\_policy** (STRING): Recommended service policy.

## vifmgr.svcPolicy.remSvc

### Severity

NOTICE

### Description

This message occurs when a service is removed from a network interface service policy.

### Corrective Action

(None).

### Syslog Message

Service %s has been removed from network interface service policy %s on Vserver %s.

### Parameters

**service** (STRING): Name of the service that has been removed.

**policy** (STRING): Name of the affected service policy.

**vserver** (STRING): Name of the Vserver associated with this policy.

## vifmgr.unabletohostmovedlif events

### vifmgr.unabletohostmovedlif

### Severity

NOTICE

### Description

This message occurs when an attempt to move a logical interface (LIF) IP address to a specific node/port failed, and a subsequent attempt was made to move that LIF to another location.

### Corrective Action

(None).

### Syslog Message

LIF %s (on virtual server %s), IP address %s, could not be hosted and is being moved to node %s, port %s (error %s).

### Parameters

**vifname** (STRING): Name of the LIF that is being moved.

**vsrvid** (STRING): Name of the virtual server identifier on which this LIF resides.

**vifip** (STRING): The logical interface's IP address that is being moved.

**tonode** (STRING): Name of the node to where the LIF is moving.

**toport** (STRING): Name of the port to where the LIF is moving.

**error** (STRING): Error that was reported.

## vifmgr.viplifwithoutbgp events

### vifmgr.vipLifWithoutBgp

**Severity**

ERROR

**Description**

This message occurs when a logical interface (LIF) associated with the Virtual IP (VIP) feature does not have a BGP peer group.

**Corrective Action**

Use the "network bgp peer-group create" command to create a BGP peer group with the same IPspace and node as the affected LIF.

**Syslog Message**

VIP lif %s on Vserver %s does not have BGP peer group in IPspace %s on node %s.

**Parameters**

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

**ipspace\_name** (STRING): IPspace name.

**node\_name** (STRING): Name of the node to which LIFs could not be replicated.

## vifmgr.vs events

### vifmgr.vs.cache.updateFailed

**Severity**

ALERT

**Description**

This message occurs when Logical Interface Manager (VIFMgr) is unable to retrieve the list of Vservers from the management gateway Vserver cache in the allotted time. VIFMgr attempts to recover from this issue. If unsuccessful, clients might not be able to access data through this node.

**Corrective Action**

If the problem persists, contact Contact NetApp technical support. for assistance.

**Syslog Message**

VIFMgr is unable to receive a Vserver cache update from the management gateway.

**Parameters**

(None).

# virtual events

## virtual.limits events

### virtual.limits.disabled

#### Severity

INFORMATIONAL

#### Description

This message occurs when the process virtual memory limit calculations are disabled on a given node. The virtual memory calculations are explicitly disabled by setting the boot argument "bootarg.engage\_vss\_limit=false".

#### Corrective Action

(None).

#### Syslog Message

Virtual memory limit calculations are disabled on node %s.

#### Parameters

**node** (STRING): Name of the node where the calculations are disabled.

### virtual.limits.proc

#### Severity

INFORMATIONAL

#### Description

This message occurs when the process virtual limit is overridden by the /mroot/etc/cluster\_config/vsz\_override file. An entry in this file indicates that the default virtual size of a process needed to be overridden.

#### Corrective Action

(None).

#### Syslog Message

Virtual memory limit for process %s on node %s overridden to %llu.

#### Parameters

**process** (STRING): Process that required the override  
**node** (STRING): Node name where the file was updated.  
**size** (LONGINT): New virtual size, in bytes.



# vldb events

## vldb.adtcons events

### vldb.adtcons.rpcfailure

#### Severity

ERROR

#### Description

This message occurs when the Remote Procedure Call (RPC) from the Volume Location Database (VLDB) to the Management Gateway (MGWD) fails. This RPC updates the audit resource database when a volume belonging to a Vserver on a specific aggregate is deleted or moved from that aggregate and that volume is the last volume belonging to that Vserver residing on that aggregate.

#### Corrective Action

Update the audit resource database by creating a temporary volume belonging to the applicable Vserver and aggregate. The temporary volume can take the defaults for the following commands and does not need to be exported to clients for data access: 'volume create -volume volume\_name -aggregate aggr\_name -vserver vserverid' After creating the temporary volume, delete it manually by using the following command: 'volume delete -volume volume\_name -vserver vserverid' This process updates the audit resource database in the cluster and recovers from this error condition. If the problem persists, contact NetApp technical support for assistance.

#### Syslog Message

RPC failure linked with the consolidation job for auditing on Vserver ID '%d', and aggregate '%s'.

#### Parameters

**vserverid** (INT): ID of the Vserver that owned the last volume deleted or moved from the aggregate.  
**aggruuid** (STRING): UUID of the aggregate that owned the moved or deleted volume.

## vldb.aggrbladeid events

### vldb.aggrBladeID.missing

#### Severity

NOTICE

#### Description

This message occurs when the system discovers an inconsistency in the volume location database (VLDB). The VLDB contains information about volumes that are owned by a node with an unknown unique universal ID (UUID). This condition might arise if an aggregate is not owned by any of the nodes in the cluster, as in the case of aggregates at a MetroCluster® destination site.

#### Corrective Action

(None).

#### Syslog Message

The volume '%s' is located on the aggregate with UUID '%s' whose owning dblade UUID '%s' does not exist in the Volume Location Database.

## Parameters

**volume** (STRING): Name of the volume.  
**aggregateUUID** (STRING): UUID of the aggregate.  
**bladeUUID** (STRING): UUID of the owning node.

# vldb.aggregateName events

## vldb.aggregateName.mismatch

### Severity

INFORMATIONAL

### Description

This message occurs when an aggregate's name in the Volume Location Database (VLDB) does not match the aggregate's current name on the node. This could happen if the name of the aggregate was changed or if another aggregate in the cluster has its UUID changed to another aggregate's UUID in the cluster, as reported to the VLDB.

### Corrective Action

(None).

### Syslog Message

The aggregate '%s' with name '%s' on node '%s' does not match with the name '%s' in the VLDB.

## Parameters

**aggregate** (STRING): UUID of the aggregate.  
**newName** (STRING): New name of the aggregate on the node.  
**node** (STRING): Name of the node.  
**existingName** (STRING): Existing name of the aggregate in the VLDB.

# vldb.aggregateuuid events

## vldb.aggregateUUID.mismatch

### Severity

INFORMATIONAL

### Description

This message occurs when an aggregate's Universally Unique Identifier (UUID) in the Volume Location Database does not match the aggregate's current UUID. This could happen if the UUID of the aggregate was changed or if another aggregate in the cluster with the same name is plugged in.

### Corrective Action

(None).

### Syslog Message

The aggregate '%s' with UUID '%s' on node '%s' was changed to '%s' in the Volume Location Database.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**existingUUID** (STRING): Existing Universally Unique Identifier (UUID) of the aggregate.

**node** (STRING): Name of the node.

**newUUID** (STRING): New Universally Unique Identifier (UUID) of the aggregate.

## vldb.aggregateUUID.missing

### Severity

ERROR

### Description

This message occurs due to a Volume Location Database inconsistency. The Volume Location Database contains a volume that exists on a aggregate whose universal unique identifier (UUID) is not in the Volume Location Database. This might happen if an aggregate that contains volumes is deleted: for example, through the maintenance menu.

### Corrective Action

Contact technical support.

### Syslog Message

The volume '%s' is located on the aggregate with UUID '%s' which does not exist in the Volume Location Database.

## Parameters

**volume** (STRING): Name of the volume

**aggregateUUID** (STRING): Universal Unique Identifier (UUID) of the aggregate

## vldb.aggrupdate events

### vldb.aggrUpdate.noRepl

### Severity

ERROR

### Description

This message occurs in a MetroCluster(tm) configuration, when the blade UUID (Universal Unique Identifier) is updated on the local cluster, but a failure prevents the aggregate ownership update from being replicated to the remote cluster.

### Corrective Action

Contact NetApp technical support for assistance.

### Syslog Message

The ownership update of aggregates on blade UUID '%s' was not replicated to the remote cluster.

## Parameters

**bladeUUID** (STRING): UUID of the Dblade owning the aggregate.

# vldb.aggrvote events

## vldb.aggrVote.noRepl

### Severity

INFORMATIONAL

### Description

This message occurs when a row of the aggr vote table is not replicated to the local table when an update is received from the remote cluster. Replication is skipped if the storage and NVRAM mirroring connectivity was indicated as down when the record was updated.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**aggrUUID** (STRING): The UUID of the aggregate whose mirror vote was not replicated.

**localNplexes** (INT): The value of nplexes in the local table.

**localPlex0** (INT): The value of plex0 id in the local table.

**localPlex1** (INT): The value of plex1 id in the local table.

**remoteNplexes** (INT): The value of nplexes from the remote cluster update.

**remotePlex0** (INT): The value of plex0 id from the remote cluster update.

**remotePlex1** (INT): The value of plex1 id from the remote cluster update.

# vldb.mccaggregate events

## vldb.mccaggregate.rename

### Severity

INFORMATIONAL

### Description

This message occurs during MetroCluster(tm) configuration setup, when the aggregate at the destination site is renamed to handle a name collision. This can happen when an aggregate is renamed at the source site, which causes its mirror to be renamed at the destination site. It can also happen when an aggregate is created at the source site, which causes its mirror to be created at the destination site. In both cases, the issue is that an aggregate of the same name already exists on the destination site, resulting in a name collision. The aggregate at the destination is renamed to resolve the name collision.

### Corrective Action

If desired, rename the aggregate at the source site to a name that does not collide with an existing name at the destination site. You can choose any new name that does not cause a name collision; the new name will be propagated to the destination site. If the old name is used in any scripts, you should also change those to the new name. If you choose not to rename the aggregate at the source site, the aggregate will have different names at the source and destination sites, but functionality is not affected.

## Syslog Message

The aggregate with UUID '%s' and name '%s' is renamed at the MetroCluster configuration destination site as '%s'.

## Parameters

**uuid** (STRING): Universally Unique Identifier (UUID) of the aggregate.

**aggrNameStr** (STRING): Name of the aggregate at the source site.

**aggrRenameStr** (STRING): Name of the aggregate at the destination site.

# vldb.update events

## vldb.update.duringsofail

### Severity

EMERGENCY

### Description

This message occurs when the aggregate ownership update fails during switchover. The update is initiated by the management gateway daemon (mgwd), to fetch aggregate information from all of the nodes in the cluster and to update the volume location database (VLDB).

### Corrective Action

Contact NetApp technical support for assistance.

## Syslog Message

Aggregate ownership update failed for node '%s'.

## Parameters

**hostname** (STRING): Name of the node.

# vldb.vldbagggr events

## vldb.vldbAggr.missing

### Severity

ERROR

### Description

This message occurs when an aggregate is identified that does not exist in the Volume Location Database. This could happen if a disk shelf that constitutes an aggregate is plugged in or when an aggregate is created in Maintenance mode.

### Corrective Action

If this is a new aggregate that must be added to the Volume Location Database, then use the (privilege: diag) "debug vreport" command to resolve the inconsistency in the VLDB: 1. Run "debug vreport show" to show the object name. 2. Run "debug vreport fix -type aggregate -object object-name". If this is a stale aggregate that must be removed, then contact support personnel.

## Syslog Message

The aggregate "%s" with UUID "%s" on node "%s" does not exist in the Volume Location Database.

## Parameters

**aggregate** (STRING): Name of the aggregate.

**uuid** (STRING): Universally Unique Identifier (UUID) of the aggregate.

**node** (STRING): Name of the node.

# vldb.vsrootvol events

## vldb.vsRootVol.del

### Severity

INFORMATIONAL

### Description

This message occurs when the Vserver root volume is force deleted while non-root volumes for the Vserver still exist in the cluster. Under normal circumstances, the Vserver root volume cannot be deleted until all the other volumes for the Vserver are deleted. Although deleting the Vserver root volume does not change the operational state for the Vserver, volume operations that refer to the Vserver root volume for information will fail if the corrective action is not taken.

### Corrective Action

Assign an already existing data volume as the new root volume for the Vserver whose root volume was deleted by using the 'volume make-vsroot -vserver vsserver\_name -volume new\_root\_volume\_name' command.

### Syslog Message

The root volume '%s' for the Vserver with Vserver ID '%d' was force deleted while non-root volumes for the Vserver still exist in the cluster.

## Parameters

**volume** (STRING): Root volume name that was deleted.

**vserverid** (INT): ID for the Vserver that owned the deleted root volume.

# vm events

## vm.dev events

### vm.dev.diskSlowIO

#### Severity

ERROR

#### Description

This message occurs when Data ONTAP® is running on a virtual platform, and slow I/O is detected.

#### Corrective Action

Collect virtual platform logs and trigger a performance archive AutoSupport® message with the "autosupport invoke-performance-archive" command, and then contact NetApp technical support for troubleshooting assistance.

#### Syslog Message

The %s request on disk %s took %u seconds (data size %u bytes, threshold %u seconds).

#### Parameters

**op** (STRING): Read/write operation.

**disk** (STRING): Disk name.

**secUsed** (INT): Seconds the I/O spent.

**size** (INT): Data size.

**secLimit** (INT): Threshold in seconds that triggers the warning.

## vm.dev.writeprotected

#### Severity

ERROR

#### Description

This message occurs when ONTAP® software is running on a virtual platform, and a write-protected device is detected. The device is not presented to ONTAP.

#### Corrective Action

Detach the device, and then reattach the device as writable.

#### Syslog Message

The device %s is write-protected and is not presented to ONTAP.

#### Parameters

**device** (STRING): Device name.

## vm.host events

## **vm.host.error**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a VM is paused by the hypervisor. The backing datastore might have been filled or connectivity to the datastore might have been lost. Corrective actions should be performed as soon as possible.

### **Corrective Action**

Check the hypervisor logs to determine whether the backing datastore is full or connectivity to the datastore was lost, and then take the appropriate corrective action.

### **Syslog Message**

The VM is paused possibly due to a full datastore or connectivity to the datastore might have been lost. Check the hypervisor logs for errors, and then perform any necessary corrective actions as soon as possible.

### **Parameters**

(None).



# vnm events

## vnm.dataip events

### vnm.dataip.added

#### Severity

NOTICE

#### Description

This message occurs when a data IP address is being added to a node. Data access through this IP address is now available.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s is being added to a node %s (ID: %u, Uuid: %s).

#### Parameters

**dataip** (STRING): The IP address that is being added.

**nodename** (STRING): Name of the node to which the data IP is being added.

**nodeid** (INT): Unique identifier of the node to which the data IP is being added.

**nodeUuid** (STRING): Universal unique identifier of the node to which the data IP is being added.

### vnm.dataip.removed

#### Severity

NOTICE

#### Description

This message occurs when a data IP address is being removed from a node. Data access through this IP address is now unavailable.

#### Corrective Action

(None).

#### Syslog Message

Data IP address %s is being removed from node %s (ID: %u, Uuid: %s).

#### Parameters

**dataip** (STRING): The IP address that is being removed.

**nodename** (STRING): Name of the node from which the data IP is being removed.

**nodeid** (INT): Unique identifier of the node from which the data IP is being removed.

**nodeUuid** (STRING): Universal unique identifier of the node to which the data IP is being removed.

# vnvram events

## vnvram.dma events

### vnvram.dma.long.wait

**Severity**

ALERT

**Description**

This message occurs when a vNVRAM flush operation on an ONTAP-v system is taking too much time. This can indicate back-end storage issues or an internal error. The system might also reboot.

**Corrective Action**

If the problem persists, check for back-end storage issues with the system disk.

**Syslog Message**

vNVRAM flush taking over %d seconds.

**Parameters**

`vnvr_wait_sec` (INT): Time in seconds for the vNVRAM flush operation.

## vnvram.flush events

### vnvram.flush.ems

**Severity**

NOTICE

**Description**

This message occurs when the write to the backing storage device for the EMS region of VNV RAM (virtual NVRAM) takes longer than 500 msec.

**Corrective Action**

(None).

**Syslog Message**

EMS region write to VNV RAM persistent storage took %d seconds.

**Parameters**

`vnvr_write_time` (LONGINT): Time in seconds taken for the EMS region write to persistent storage.

## vnvram.flushed events

### vnvram.flushed

**Severity**

NOTICE

**Description**

This message occurs when the VNVRAM contents of any ONTAP-based virtual appliance that is not using full-mode VNVRAM, are written to the backing storage device. The writes would have occurred during either a clean or dirty shutdown.

**Corrective Action**

(None).

**Syslog Message**

Write to VNVRAM persistent storage is complete (%s).

**Parameters**

**vnvr\_reg** (STRING): Section of NVRAM that was written to the backing device.

# vol events

## vol.clone events

### vol.clone.split.optimized

#### Severity

NOTICE

#### Description

This message occurs when the clone split of a volume clone takes the optimized path which uses an Azure copy blob operation to copy objects from the parent volume to the clone.

#### Corrective Action

(None).

#### Syslog Message

Clone split of volume "%s" hosted on FabricPool "%s" takes the optimized path using Azure copy blob operation.

#### Parameters

**volume\_name** (STRING): Volume clone.

**aggregate\_name** (STRING): Aggregate which hosts the volume clone.

## vol.duplicate events

### vol.duplicate.msids

#### Severity

EMERGENCY

#### Description

This message occurs when a duplicate master set ID (MSID) has been detected. This situation might cause ONTAP to incorrectly group together the constituents of separate FlexGroup volumes. This might result in inconsistencies, as ONTAP creates directories and files across separate FlexGroup volumes. This might lead ONTAP to mark FlexGroup volume constituents as inconsistent and take them offline. ONTAP might even mark the entire aggregate as inconsistent and take it offline.

#### Corrective Action

Contact NetApp technical support immediately.

#### Syslog Message

Duplicate MSID "%lu" has been detected in SVM "%s". The duplicate MSID is used for volumes "%s" and "%s".

#### Parameters

**volumeMsid** (LONGINT): MSID of the volumes.

**vserverName** (STRING): Name of the storage virtual machine (SVM) on which the volumes reside.

**firstVolName** (STRING): Name of the first volume with the duplicate MSID.

**secondVolName** (STRING): Name of the second volume with the duplicate MSID.

# vol.inconsistent events

## vol.inconsistent.modify

### Severity

INFORMATIONAL

### Description

This message occurs when the "vol inconsistent" command is used to modify a volume's state.

### Corrective Action

(None).

### Syslog Message

Volume %s: 'vol inconsistent' command used to change inconsistent state to '%s'.

### Parameters

**vol** (STRING): Name of the volume.

**state** (STRING): State of the action "set" or "clear".

# vol.language events

## vol.language.changed

### Severity

INFORMATIONAL

### Description

This event is issued when the language of an aggregate or a volume is changed.

### Corrective Action

(None).

### Syslog Message

Language on %s %s changed to %s

### Parameters

**vol\_type** (STRING): Whether this event is generated for an aggregate or a flexible volume.

**vol** (STRING): The name of the volume or aggregate

**lang** (STRING): The name of the changed language of the volume or aggregate.

## vol.language.openFailed

### Severity

ERROR

### Description

This message occurs when an attempt to open the language file fails.

### Corrective Action

Check whether the language file exists. Retry the command based on the reason for failure.

### Syslog Message

Failed to open the language file %s, error message: %s

### Parameters

**file\_name\_location** (STRING): Language file name and location.

**error\_str** (STRING): Reason for the error.

## vol.language.unspecified

### Severity

INFORMATIONAL

### Description

This event is issued when ONTAP detects that there is no language specified for the volume or aggregate. The referenced volume or aggregate language will default to "C" (POSIX). See the na\_vol.1 man page for the list of supported languages.

### Corrective Action

If the volume is not read-only, use the 'vol lang' command to correct the language of the volume. See the na\_vol.1 man page for details about using the 'vol lang' command.

### Syslog Message

Language not set on %s %s. Using language config "C". %s

### Parameters

**vol\_type** (STRING): Whether this error occurred for an aggregate or a flexible volume

**vol** (STRING): The name of the volume or aggregate

**lang\_command** (STRING): If the volume is not read-only, this parameter contains an explanation of what you can do to set the language for the volume.

## vol.log events

### vol.log.overalloc

### Severity

ALERT

### Description

This message occurs when the volume or aggregate allocates more space than it can honor by way of reservations, or the aggregate has allocated more space than it can honor by way of guarantees. If the reserved or guaranteed space is consumed, there is insufficient physical space, which can cause the volume or aggregate to be taken offline.

### Corrective Action

Create space by increasing the volume or aggregate size, deleting data, deleting Snapshot® copies, or changing the provisioning from thick to thin. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot® copies, use the "volume snapshot delete" command. To change provisioning in a volume, reserved files can be unreserved by using the "volume file reservation" command.

To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot® copies are deleted automatically when the aggregate is full. To change provisioning of a volume in an aggregate, change the volume guarantee from "volume" to "none" by using the "space-guarantee" field of the "volume modify" command.

### Syslog Message

%s %s%s%s is logically overallocated, using %s%% logical space and %s%% physical space.

### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).

**name** (STRING): Name of this object.

**app** (STRING): Application UUID.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the object's Vserver, if the object is a volume. Otherwise, this string is empty.

**percent\_logically\_full** (STRING): Used capacity of the logical space of the object, as a percent. This value can be greater than 100 due to space reserved for metadata. The percentage is computed for the entire object, including Snapshot® copy space, and is equivalent to the "Total Used" field in the output of the "volume show-space" or "aggregate show-space" commands.

**percent\_physically\_full** (STRING): Used capacity of the physical space of the object, as a percent. This value can be greater than 100 due to space reserved for metadata. The percentage is computed for the entire object, including Snapshot® copy space, and is equivalent to the "Total Physical Used" field in the output of the "volume show-space" or "aggregate show-space" commands.

## vol.log.overalloc.inc.sav

### Severity

ALERT

### Description

This message occurs when the volume allocates more space than it can honor. Allocated space is based on all of the physical used blocks, reservations, and space saved by storage efficiency features.

### Corrective Action

Create space by increasing the volume size, deleting data or their reservations, or deleting Snapshot copies. To increase volume size, use the "volume size" command. To delete Snapshot copies, use the "volume snapshot delete" command.

### Syslog Message

%s %s%s%s (UUID %s) is logically overallocated using %s%% space, which considers space saved by storage efficiency features as well as physically used blocks and reservations.

### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies. This event should be applicable only to volumes with logical space enforcement enabled.

**name** (STRING): Name of this object.

**app** (STRING): Application Universally Unique ID (UUID).

**vserver\_uuid** (STRING): UUID of the object's or volume's Vserver.

**object\_uuid** (STRING): UUID of the object or volume.

**percent\_full** (STRING): Used capacity of the logical space of the object, along with reservations and space saved by the storage efficiency features, as a percentage. This value can be greater than 100 due to space reserved for metadata. The percentage is computed for the entire object, including Snapshot(tm) copy space, and is equivalent to the "logical-used" field in the output of the "volume show-space" command.

# vol.max events

## vol.max.size.increased

### Severity

INFORMATIONAL

### Description

This message occurs as a reminder that the maximum volume size limit is increased from the default value defined in the Platform Specific Module (PSM) table.

### Corrective Action

(None).

### Syslog Message

Maximum volume size limit is increased from %ld TB to %ld TB.

### Parameters

**old\_limit** (LONGINT): Original volume size limit as defined in the PSM table.

**new\_limit** (LONGINT): New maximum volume size limit.

# vol.mount events

## vol.mount.readOnly

### Severity

INFORMATIONAL

### Description

This message occurs when a volume or aggregate is mounted as read-only.

### Corrective Action

(None).

### Syslog Message

Mounting vol/aggr '%s' readonly due to 'online -ro' command.

### Parameters

**name** (STRING): Name of the volume or aggregate.

# vol.phy events

## vol.phy.overalloc.inc.sav

### Severity

EMERGENCY

### Description

This message occurs when the volume is running dangerously low on free blocks to write to. Free blocks are estimated based on the space saved by the storage efficiency features as well as with physically used



blocks.

### Corrective Action

Create space by increasing the volume size, deleting data, or deleting Snapshot copies. To increase volume size, use the "volume size" command. To delete Snapshot copies, use the "volume snapshot delete" command.

### Syslog Message

%s %s%s%s (UUID %s) is physically overallocated using %s%% space, which considers space saved by storage efficiency features as well as physically used blocks.

### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies. This event should be applicable only to volumes with logical space enforcement turned on.

**name** (STRING): Name of this object.

**app** (STRING): Application Universally Unique ID (UUID).

**vserver\_uuid** (STRING): UUID of the object's or volume's Vserver.

**object\_uuid** (STRING): UUID of the object or volume.

**percent\_full** (STRING): Used capacity of the physical space of the object, along with space saved by the storage efficiency features, as a percentage. This value can be greater than 100 due to efficiency features and space reserved for metadata. The percentage is computed for the entire object, including Snapshot(tm) copy space.

## vol.phys events

### vol.phys.overalloc

#### Severity

EMERGENCY

#### Description

This message occurs when the volume or aggregate is running dangerously low on free blocks to write to.

#### Corrective Action

Create space by increasing the volume or aggregate sizes, or by deleting data or deleting Snapshot® copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot® copies, use the "volume snapshot delete" command. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot® copies are deleted automatically when the aggregate is full.

#### Syslog Message

%s %s%s%s is physically overallocated, using %s%% logical space and %s%% physical space.

#### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).

**name** (STRING): Name of this object.

**app** (STRING): Application UUID.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the object's Vserver, if the object is a volume. Otherwise, this string is empty.

**percent\_logically\_full** (STRING): Used capacity of the logical space of the object, as a percent. This value can be greater than 100 due to space reserved for metadata. The percentage is computed for the entire object, including Snapshot® copy space, and is equivalent to the "Total Used" field in the output of the

"volume show-space" or "aggregate show-space" commands.

**percent\_physically\_full** (STRING): Used capacity of the physical space of the object, as a percent. This value can be greater than 100 due to space reserved for metadata. The percentage is computed for the entire object, including Snapshot® copy space, and is equivalent to the "Total Physical Used" field in the output of the "volume show-space" or "aggregate show-space" commands.

## vol.registryget events

### vol.registryGet.no512520BPS

#### Severity

NOTICE

#### Description

This message occurs when a registry request finds no 512 or 520 BPS (bytes per sector) disks.

#### Corrective Action

(None).

#### Syslog Message

No disk has 512 or 520 BPS.

#### Parameters

(None).

### vol.registryGet.nonDefBPS

#### Severity

NOTICE

#### Description

This message occurs when a registry request finds disks whose BPS (bytes per sector) is other than the default 512 or 520.

#### Corrective Action

(None).

#### Syslog Message

Disk had %d BPS; system expected 512 or 520 BPS.

#### Parameters

**cad\_bps** (INT): Disk configuration attribute of bytes per sector.

### vol.registryGet.noZoneCksum

#### Severity

NOTICE

#### Description

This message occurs when a registry request for zoned or block checksum volumes finds no such volumes.

**Corrective Action**

(None).

**Syslog Message**

No volume has zoned or block checksums.

**Parameters**

(None).

## vol.reserve events

**vol.reserve.grew****Severity**

ERROR

**Description**

This message occurs when WAFL® grows the reserve of a flexible volume that was running dangerously low on blocks for writing.

**Corrective Action**

Create space by increasing the volume size, deleting data, or changing the provisioning from thick to thin. To increase a volume's size, use the "volume size" command. To change provisioning in a volume reserved files can be unreserved by using the "volume file reservation" command.

**Syslog Message**

WAFL grew reserve of volume %s%s%s because the volume was running low on blocks.

**Parameters**

**name** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**vserver\_uuid** (STRING): Name of the associated Vserver's UUID.

# volaggr events

## volaggr.offline events

### volaggr.offline

#### Severity

ALERT

#### Description

Some aggregates are forced offline. volume creation could cause duplicate FSIDs.

#### Corrective Action

Bring all offline aggregates online.

#### Syslog Message

Some aggregates are offline. Volume creation could cause duplicate FSIDs.

#### Parameters

(None).

# volcmd events

## volcmd.snapmirror events

### volcmd.snapmirror.removed

**Severity**

ERROR

**Description**

This message occurs to advise you that the 'vol snapmirror' set of commands has been removed.

**Corrective Action**

Use the 'snapmirror' command instead.

**Syslog Message**

The 'vol snapmirror' set of commands has been removed; use the 'snapmirror' command instead.

**Parameters**

(None).

# volcopy events

## volcopy.dst events

### volcopy.dst.autoRestrictMsg

**Severity**

NOTICE

**Description**

This event is issued when a vol copy start is issued on a volume which is offline.

**Corrective Action**

Vol copy start is issued on a restricted volume, use "vol restrict" to restrict the destination volume.

**Syslog Message**

Vol copy automatically restricting destination volume %s (currently offline).

**Parameters**

**dst** (STRING): The vol copy destination volume

### volcopy.dst.metadataOnly

**Severity**

NOTICE

**Description**

This event is issued when a vol copy start operation is issued on a volume which is reverting its checksum state.

**Corrective Action**

(None).

**Syslog Message**

Metadata only vol copy from %s to %s.

**Parameters**

**src** (STRING): The vol copy source volume

**dst** (STRING): The vol copy destination volume

### volcopy.dst.noMemory

**Severity**

ERROR

**Description**

This event is issued when there is insufficient memory on the vol copy destination for the transfer to continue.

**Corrective Action**

(None).

**Syslog Message**

System is low on memory, aborting vol copy destination transfer.

**Parameters**

(None).

**volcopy.dst.revert****Severity**

ERROR

**Description**

This event is issued when a vol copy start operation has to revert the checksum state of the destination volume because the source and destination volumes have differing checksum states.

**Corrective Action**

No corrective action required.

**Syslog Message**

A vol copy transfer to destination volume %s has found that the source and destination volumes have differing checksum states. Data ONTAP is aborting the transfer and reverting the checksum state for the destination volume. After Data ONTAP displays a message showing the revert is complete, you will need to retry the operation. You can use the "vol status -c" command to see the progress of the reversion.

**Parameters**

**dst** (STRING): The vol copy destination volume

**volcopy.dst.reverting****Severity**

ERROR

**Description**

This event is issued when a vol copy start operation is issued on a volume which is reverting its checksum state.

**Corrective Action**

(None).

**Syslog Message**

Aborting vol copy transfer to volume %s because it is reverting its checksum state.

**Parameters**

**dst** (STRING): The vol copy destination volume

## volcopy.flextotraderr events

### volcopy.flexToTradErr

**Severity**

ERROR

**Description**

This message occurs when the user tries to copy a flexible volume into a traditional volume using the command 'vol copy'.

**Corrective Action**

Retry the operation with a flexible volume as the destination.

**Syslog Message**

vol copy: A flexible volume cannot be copied into a traditional volume.

**Parameters**

(None).

## volcopy.src events

### volcopy.src.noMemory

**Severity**

ERROR

**Description**

This event is issued when there is insufficient memory on the vol copy source for the transfer to continue.

**Corrective Action**

(None).

**Syslog Message**

System is low on memory, aborting vol copy source.

**Parameters**

(None).

### volcopy.src.revertPending

**Severity**

NOTICE

**Description**

This message indicates the source is being reverted by WAFL and cannot currently be used for a 'vol copy' operation.



### **Corrective Action**

This message indicates the source is being reverted. After the revert completes the operation should be retried.

### **Syslog Message**

vol copy: The source is being reverted by WAFL

### **Parameters**

(None).

## **volcopy.src.snapshotErr**

### **Severity**

ERROR

### **Description**

This event is issued when we are not able to obtain all the snapshot information, from the vol copy source volume, required for a vol copy transfer.

### **Corrective Action**

(None).

### **Syslog Message**

Vol copy not able to obtain snapshot information from source volume %s

### **Parameters**

**src** (STRING): The vol copy source volume

## **volcopy.tradtoflexerr events**

### **volcopy.tradToFlexErr**

### **Severity**

ERROR

### **Description**

This message occurs when the user tries to copy a traditional volume into a flexible volume using the command 'vol copy'.

### **Corrective Action**

Retry the operation with a traditional volume as the destination.

### **Syslog Message**

vol copy: A traditional volume cannot be copied into a flexible volume.

### **Parameters**

(None).

# volcreate events

## volcreate.fsid events

### volcreate.fsid.dup

#### Severity

ERROR

#### Description

Flexible volumes were created with file system identifiers (FSIDs) which conflict with existing volumes. This can happen when aggregates are forced offline.

#### Corrective Action

1. run 'aggr online' on all offlined aggregates. some of the old vvol will be forced offline due to the FSID conflict.
2. run 'vol restrict' and 'vol rewrite\_fsid' on the new vvol
3. run 'vol online' on each flexible volume that was forced offline in step 1 above.

#### Syslog Message

Volume %s's fsid 0x%x conflicts with a volume owned by vfiler %s.

#### Parameters

**volName** (STRING): Name of the volume created

**FSID** (INT): FSID of the volume

**vfiler\_name** (STRING): The vfiler which owns the existing offlined volume with the same FSID.

## volcreate.name events

### volcreate.name.dup

#### Severity

ERROR

#### Description

Flexible volumes were created with names which conflict with existing volume names. This can happen when aggregates are forced offline.

#### Corrective Action

if this is not the root volume of the vfiler: 1. run 'vfiler remove' on the vfiler and the new vvol 2. run 'vol rename' on the new vvol 3. resolve FSID conflicts if any 4. add the old vvol back to the vfiler if this is the root of volume of the vfiler: 1. run 'vfiler destroy' on the vfiler 2. run 'vol rename' on the new vvol 3. resolve FSID conflicts if any 4. run 'vfiler create -r' on the vfiler

#### Syslog Message

Volume '%s' has a name which conflicts with a volume owned by vfiler '%s'.

#### Parameters

**volName** (STRING): Volume being created

**vfiler\_name** (STRING): This vfiler owns an existing volume with the same name.

# volmigrate events

## volmigrate.error events

### volmigrate.error

#### Severity

INFORMATIONAL

#### Description

Description of the reason a volume migrate failed or was not attempted.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**volName** (STRING): Volume being migrated

**recipientFiler** (STRING): Name of the filer to which the volume was being migrated

**error** (STRING): The specific error string that details why migrate failed.

## volmigrate.migrating events

### volmigrate.migrating

#### Severity

INFORMATIONAL

#### Description

A volume is being migrated from the filer by means of a SAN ownership change.

#### Corrective Action

(None).

#### Syslog Message

Migrating volume %s to %s (ID %s).

#### Parameters

**volName** (STRING): Volume being migrated

**recipientFiler** (STRING): Name of the filer to which the volume is being migrated

**recipientSysID** (STRING): System ID of the filer to which the volume is being migrated

## volmigrate.result events

## **volmigrate.result**

### **Severity**

INFORMATIONAL

### **Description**

Description of the result of a volume migration.

### **Corrective Action**

(None).

### **Syslog Message**

Migration of volume %s to %s %s.

### **Parameters**

**volName** (STRING): Volume being migrated

**recipientFiler** (STRING): Name of the filer to which the volume is being migrated

**result** (STRING): A string describing the result of the migration: either 'was successful' or 'failed'.

# volume events

## volume.check events

### Volume.Check.Completed

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

NOTICE

#### Description

This message occurs when a 'volume check' scan successfully runs to completion.

#### Corrective Action

(None).

#### Syslog Message

The 'volume check' scan was completed for volume [%s] on Vserver [%s].

#### Parameters

**volume** (STRING): Name of the constituent volume.

**vserver** (STRING): Vserver hosting the volume.

### Volume.Check.Failed

#### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

#### Severity

ERROR

#### Description

This message occurs when a 'volume check' scan fails to run to completion.

#### Corrective Action

Correct the reported error and reissue the 'volume check' command. If error persists, contact NetApp technical support.

#### Syslog Message

The 'volume check' scan failed for volume [%s] on Vserver [%s] with error [%s].

#### Parameters

**volume** (STRING): Name of the constituent volume.

**vserver** (STRING): Vserver hosting the volume.

**error** (STRING): Error string.

# Volume.Check.Report

## Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

## Severity

INFORMATIONAL

## Description

This message occurs when volume check completes the scan. The report contains information about the following inconsistencies detected and the corrective action taken: 1) Corruption/loss of redirector content. 2) Valid redirector content file, but data file being pointed to did not exist in the Data Constituents. 3) Create Dungle files. 4) NLINKS inconsistency. 5) ACL mismatch. 6) Corrupt data file name. 7) Data file does not have a corresponding redirector file. 8) Corruption/loss of stream redirector content. 9) Valid redirector stream content, but stream data being pointed to did not exist in the Data Constituents. 10) Corrupt stream data file name. 11) Stream data does not have a corresponding redirector stream.

## Corrective Action

(None).

## Syslog Message

Volume Check Report for Constituent Volume [%s] in Vserver [%s] - Inconsistency - (Issues Found/Issues Fixed); Corrupt Redirector - (%d/%d); Dangling Redirector - (%d/%d); Temporary Files - (%d/%d); NLinks - (%d/%d); ACL mismatch - (%d/%d); Dangling Data File - (%d/%d); Corrupt Data File - (%d/%d); Corrupt Stream Redirector - (%d/%d); Dangling Stream Redirector - (%d/%d); Dangling Stream Data - (%d/%d); Corrupt Stream Data - (%d/%d);

## Parameters

**volume** (STRING): Constituent volume.  
**vserver** (STRING): Vserver hosting the volume.  
**numOfCorruptRedirectorsFound** (INT): Number of corrupt or lost redirectors found.  
**numOfCorruptRedirectorsFixed** (INT): Number of corrupt or lost redirectors fixed.  
**numOfDanglingRedirectorsFound** (INT): Number of dangling redirectors found.  
**numOfDanglingRedirectorsFixed** (INT): Number of dangling redirectors fixed.  
**numOfTempFilesFound** (INT): Number of hidden temporary files remaining after failed creates.  
**numOfTempFilesRemoved** (INT): Number of hidden temporary files removed.  
**numOfNlinkMisMatchesFound** (INT): Number of NLINKS mismatches found.  
**numOfNlinkMisMatchesFixed** (INT): Number of NLINKS mismatches fixed.  
**numOfACLMisMatchesFound** (INT): Number of ACL mismatches found.  
**numOfACLMisMatchesFixed** (INT): Number of ACL mismatches fixed.  
**numOfDanglingDataFilesFound** (INT): Number of dangling data files found.  
**numOfDanglingDataFilesFixed** (INT): Number of dangling data files fixed.  
**numOfCorruptDataFilesFound** (INT): Number of corrupt or lost data files found.  
**numOfCorruptDataFilesFixed** (INT): Number of corrupt or lost data files fixed.  
**numOfCorruptStreamRedirectorsFound** (INT): Number of corrupt stream redirectors found.  
**numOfCorruptStreamRedirectorsFixed** (INT): Number of corrupt stream redirectors fixed.  
**numOfDanglingStreamRedirectorsFound** (INT): Number of dangling stream redirectors found.  
**numOfDanglingStreamRedirectorsFixed** (INT): Number of dangling stream redirectors fixed.  
**numOfDanglingStreamDataFound** (INT): Number of dangling stream data instances found.  
**numOfDanglingStreamDataFixed** (INT): Number of dangling stream data instances fixed.  
**numOfCorruptStreamDataFound** (INT): Number of corrupt stream data instances found.  
**numOfCorruptStreamDataFixed** (INT): Number of corrupt stream data instances fixed.

## Volume.Check.Start.Job.Started

### Deprecated

Deprecated as of Data ONTAP 9.6 because Infinite Volumes no longer exist.

### Severity

INFORMATIONAL

### Description

This message occurs when the 'volume check start' command starts the job, but fails to update the RDB(replicated database) table and the corresponding job is still running.

### Corrective Action

Use the 'job stop' command to abort the reported job before starting another check.

### Syslog Message

A 'volume check' command started job [ID: %s] for volume [%s] in Vserver [%s].

### Parameters

**jobid** (STRING): Job ID of the scan.

**volume** (STRING): Name of the constituent volume for which the job was started.

**vserver** (STRING): Vserver hosting the volume.

## volume.transition events

### volume.transition.plugin.slow

#### Severity

INFORMATIONAL

#### Description

This message occurs when a volume transition plug-in is taking longer than expected.

#### Corrective Action

(None).

#### Syslog Message

Volume transition on volume %s, plug-in %s running %d ms (expected %d ms).

#### Parameters

**voluuid** (STRING): Unique identifier for this volume.

**plugin** (STRING): Name of the slow plug-in.

**used** (INT): Runtime for the plug-in, in milliseconds.

**expected** (INT): Expected runtime for the plug-in, in milliseconds.

# vpeer events

## vpeer.name events

### vpeer.name.conflict

#### Severity

NOTICE

#### Description

This message occurs when system generated a new name for a Vserver because of a naming conflict with another Vserver in the local cluster.

#### Corrective Action

(None).

#### Syslog Message

Peer Vserver name conflict in local cluster with existing local or peer Vserver %s. Name conflict solved using system generated name %s.

#### Parameters

**peer\_vserver** (STRING): Vserver name in conflict.

**unique\_peer\_vserver** (STRING): Replacement Vserver name generated by the system to resolve the name conflict.

## vpeer.request events

### vpeer.request.diff.ipspace

#### Severity

NOTICE

#### Description

This message occurs when the Vserver peer request is received and Vserver peer permission exists for the Vserver and peer-cluster, but the Vserver belongs to an ipspace that is different from the cluster peer relationship ipspace. The Vserver peer request will remain in the pending state unless explicitly accepted.

#### Corrective Action

Use the "vserver peer accept" command in the destination cluster to explicitly accept the Vserver peer request, and use the "vserver peer show" command to verify that the Vserver peer goes to peered state.

#### Syslog Message

"Vserver peer request for the peer-Vserver '%s' and the peer-cluster '%s' needs to be explicitly accepted, because the local Vserver '%s' ipspace: '%s' is different from the cluster peer relationship ipspace: '%s'."

#### Parameters

**peer\_vserver** (STRING): Peer Vserver name.

**peer\_cluster** (STRING): Peer cluster name.

**local\_vserver** (STRING): Local Vserver name.

**vserver\_ipspace** (STRING): Ipspace of the local Vserver.



**cluster\_peer\_ipspace** (STRING): Ipspace of the cluster peer relationship.

# vs events

## vs.migrate events

### vs.migrate.cleanup.completed

**Severity**

NOTICE

**Description**

This message occurs when the source Vserver cleanup operation has completed.

**Corrective Action**

(None).

**Syslog Message**

Migrate source Vserver cleanup completed for Vserver: %s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

### vs.migrate.complete

**Severity**

NOTICE

**Description**

This message occurs when migrate operation completed.

**Corrective Action**

(None).

**Syslog Message**

Migrate of Vserver:%s has completed.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

### vs.migrate.cutover.commit.started

**Severity**

NOTICE

**Description**

This message occurs when migrate cutover commit started successfully.

**Corrective Action**

(None).

**Syslog Message**

Migrate cutover has started for Vserver:%s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.cutover.completed****Severity**

NOTICE

**Description**

This message occurs when migrate cutover operation completed.

**Corrective Action**

(None).

**Syslog Message**

Migrate cutover has completed for Vserver:%s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.cutover.precommit.started****Severity**

NOTICE

**Description**

This message occurs when the migrate operation has started the precommit for cutover.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.cutover.triggered****Severity**

NOTICE

**Description**

This message occurs when cutover is triggered on a Vserver that is migrating.

**Corrective Action**

(None).

**Syslog Message**

Cutover was triggered on a migrating Vserver:%s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.dst.ponr.updated****Severity**

NOTICE

**Description**

This message occurs when the vserver migrate cutover workflow has reached and updated the point of no return (PONR) on the destination cluster.

**Corrective Action**

(None).

**Syslog Message**

Point of no return on destination Vserver %s has been updated successfully.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.failed****Severity**

ERROR

**Description**

This message occurs when the Vserver migrate operation fails. This can occur before/during cutover or during post cutover activity. If failure is before/during cutover then source Vserver is still operational and if migrate fails during post cutover activity then destination vserver is up. Once migrate issues are fixed, migrate can be resumed back to its original state.

**Corrective Action**

Run the "vserver migrate show -fields status-details" to view the error and follow the corrective action recommended.

**Syslog Message**

Migration of Vserver "%s" Failed.

## Parameters

**vserver\_name** (STRING): Vserver name.

**err\_status** (STRING): Reason for the migrate operation failure and the corrective action to fix the problem.

**migration\_uuid** (STRING): Migration UUID.

## vs.migrate.pause.complete

### Severity

NOTICE

### Description

This message occurs when the migrate operation for a Vserver is paused.

### Corrective Action

(None).

### Syslog Message

Migrate of Vserver:%s was paused.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## vs.migrate.precheck.completed

### Severity

NOTICE

### Description

This message occurs when the migrate operation has completed the prechecks.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**sourceCluster** (STRING): Name of the source cluster.

**destinationCluster** (STRING): Name of the destination cluster.

## vs.migrate.precheck.failed

### Severity

ERROR

### Description

This message occurs when the migrate operation has failed the prechecks.

### Corrective Action

Run the "vserver migrate show" command to view the description of the error and the corrective action.

### Syslog Message

(None).

### Parameters

**vserver\_name** (STRING): Name of the Vserver.

**err\_status** (STRING): Reason for the precheck operation failure and the corrective action to fix the problem.

**sourceCluster** (STRING): Name of the source cluster.

**destinationCluster** (STRING): Name of the destination cluster.

## vs.migrate.precheck.started

### Severity

NOTICE

### Description

This message occurs when the migrate operation has started the prechecks.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**vserver\_name** (STRING): Name of the Vserver.

**sourceCluster** (STRING): Name of the source cluster.

**destinationCluster** (STRING): Name of the destination cluster.

## vs.migrate.readyforcutover

### Severity

NOTICE

### Description

This message occurs when the migrate operation has reached the ready-for-cutover state.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## **vs.migrate.resume**

### **Severity**

NOTICE

### **Description**

This message occurs when the migrate operation for a Vserver is resumed after an abort or a failure.

### **Corrective Action**

(None).

### **Syslog Message**

Migrate of Vserver:%s was resumed.

### **Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## **vs.migrate.rolledback**

### **Severity**

NOTICE

### **Description**

This message occurs when Vserver migrate operation moves back from ready-to-cutover to transferring state after cutover was triggered. This can be due to high transfer duration or due to some failure that might have occurred when cutover was triggered on destination cluster.

### **Corrective Action**

(None).

### **Syslog Message**

Migrate state for Vserver %s was rolled back to transferring. Run "vserver migrate show -fields status-details" to see the reason for rollback. If the rollback is due to high transfer duration, reduce the I/O operations at the migrate source cluster to a lower rate.

### **Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## **vs.migrate.src.abort.completed**

### **Severity**

NOTICE

### **Description**

This message occurs when the Vserver migrate abort operation on the source cluster has completed.

### **Corrective Action**

(None).

## Syslog Message

(None).

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## vs.migrate.src.abort.failed

### Severity

ERROR

### Description

This message occurs when the Vserver abort operation on the source cluster has failed.

### Corrective Action

Run the "vserver migrate show -fields status-details" command to view the description of the error and the corrective action.

## Syslog Message

Vserver migrate abort operation failed for Vserver: %s. Error: %s

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**err\_status** (STRING): Reason for the abort operation failure and the corrective action to fix the problem.

**migration\_uuid** (STRING): Migration UUID.

## vs.migrate.src.cleanup.failed

### Severity

ERROR

### Description

This message occurs when the source Vserver cleanup operation fails.

### Corrective Action

Run the "vserver migrate show -fields status-details" command to view the description of the error and the corrective action.

## Syslog Message

Migrate source Vserver cleanup failed for Vserver: %s.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**err\_status** (STRING): Reason for the source cleanup operation failure and the corrective action to fix the problem.

**migration\_uuid** (STRING): Migration UUID.

## vs.migrate.src.cleanup.started



**Severity**

NOTICE

**Description**

This message occurs when the source Vserver cleanup operation starts. This operation starts only after the Vserver user data and configuration metadata have been replicated to the destination cluster, and the destination Vserver is serving data. This cleanup operation eventually deletes the source Vserver.

**Corrective Action**

(None).

**Syslog Message**

Migrate source Vserver cleanup started for Vserver: %s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.src.ponr.failed****Severity**

ERROR

**Description**

This message occurs when the vserver migrate cutover workflow has failed to reach and update the point of no return (PONR).

**Corrective Action**

Run "vserver migrate show -instance" at the source cluster. If the point of no return is set to true, run "vserver migrate restart" command at the migrate destination cluster to start the destination Vserver. If the point of no return is set to false and the source Vserver is in a stopped state, run "Vserver start" at the source cluster to start the Vserver.

**Syslog Message**

Failed to update point of no return for Vserver %s.

**Parameters**

**vserver\_name** (STRING): Vserver Name.

**err\_status** (STRING): Reason for the ponr operation failure and the corrective action to fix the problem.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.src.ponr.updated****Severity**

NOTICE

**Description**

This message occurs when the vserver migrate cutover workflow has reached and updated the point of no return (PONR) on the source cluster.

**Corrective Action**

(None).

**Syslog Message**

Point of no return on source Vserver %s has been updated successfully.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.start****Severity**

NOTICE

**Description**

This message occurs when the migrate operation for Vserver from source cluster to destination cluster starts.

**Corrective Action**

(None).

**Syslog Message**

Migrate of Vserver:%s started from source cluster:%s to destination cluster:%s.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**sourceCluster** (STRING): Name of the source cluster.

**destinationCluster** (STRING): Name of the destination cluster.

**migration\_uuid** (STRING): Migration UUID.

**vs.migrate.sync.reached****Severity**

NOTICE

**Description**

This message occurs when the migrate operation has reached sync.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**migration\_uuid** (STRING): Migration UUID.

## **vs.migrate.transferring.completed**

### **Severity**

NOTICE

### **Description**

This message occurs when the migrate operation has completed the transferring operation.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**vserver\_name** (STRING): Name of the Vserver.  
**migration\_uuid** (STRING): Migration UUID.

## **vs.migrate.transferring.started**

### **Severity**

NOTICE

### **Description**

This message occurs when the migrate operation has started transferring successfully.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**vserver\_name** (STRING): Name of the Vserver.  
**migration\_uuid** (STRING): Migration UUID.

## **vs.multiple events**

### **vs.multiple.adtjobs**

### **Severity**

ERROR

### **Description**

This message occurs when more than one Vserver audit consolidation jobs is found.

### **Corrective Action**

Check Vserver audit configuration. If enabled, disable it using the following commands: cluster::> vserver audit show cluster::> vserver audit disable -vserver Check for Vserver audit consolidation jobs. If still present, delete the jobs using the following commands: cluster::> job show cluster::> job expunge -id Check for consolidation processes. If any still exist, kill the processes using the following commands from

systemshell: > ps -auxw | grep adtcons > sudo kill -9 "pid of the process" Enable the Vserver audit configuration by using the following command: cluster::> vserver audit enable

### **Syslog Message**

More than one Vserver audit consolidation jobs found for Vserver %s.

### **Parameters**

**vserver** (STRING): Name of the Vserver that has more than one consolidation job.

# vsa events

## vsa.azure events

### vsa.azure.nlb.probeActive

#### Severity

NOTICE

#### Description

This message occurs when ONTAP® Cloud on Microsoft® Azure® receives a probe, indicating that a transition to the expected normal state has occurred. The port has gone to active state.

#### Corrective Action

(None).

#### Syslog Message

Received Load Balancer probe (%s) for %d ports (port range: %d to %d).

#### Parameters

**detail** (STRING): Azure Load Balancer probes were received on the ONTAP Network Load Balancer listener port.

**activated\_ports** (INT): Number of ports that received the probe.

**first\_activated\_port** (INT): First (lowest) port that received a probe.

**last\_activated\_port** (INT): Last (highest) port that received a probe.

### vsa.azure.nlb.probeInactive

#### Severity

ALERT

#### Description

This message occurs when ONTAP® Cloud on Microsoft® Azure® does not receive expected health probes from the Network Load Balancer service within the expected time limit. This might cause Azure to fail over traffic from the corresponding address. If Azure Load Balancer does not have an active address, all traffic might be blocked from reaching the configured address.

#### Corrective Action

Verify the configuration of the Azure Load Balancer and the ONTAP Network Load Balancer, as well as the status of the node management LIF. To show the state of the ONTAP Network Load Balancer, use the (privilege: diagnostic) "virtual-machine show-network-load-balancer" command. To show the inactive probes of the ONTAP Network Load Balancer, use the (privilege: diagnostic) "virtual-machine show-network-load-balancer -active false" command. In the above command check if the "Last Probe Time" is being regularly updated (approximately every 5 seconds) and the Cloud Load Balancer probe port matches the "Probe Port".

#### Syslog Message

Failed to receive Load Balancer probe (%s) for %d ports (port range: %d to %d), within %d seconds.

## Parameters

**detail** (STRING): Azure Load Balancer Probes were not received in the expected timeframe on the ONTAP Network Load Balancer listener port.  
**missed\_ports** (INT): Number of ports that did not receive a probe.  
**first\_missed\_port** (INT): First (lowest) port that did not receive a probe.  
**last\_missed\_port** (INT): Last (highest) port that did not receive a probe.  
**timeout\_in\_seconds** (LONGINT): Seconds before timeout occurs.

# vsa.cloud events

## vsa.cloud.aws.api.rtUpdt

### Severity

NOTICE

### Description

This message occurs when cluster LIF route update succeeds upon node switchover.

### Corrective Action

(None).

### Syslog Message

Command "%s" is successfully executed.

## Parameters

**action** (STRING): Amazon Web Services (AWS) API command for route table updates or description.

## vsa.cloud.aws.api.rtUpdtFail

### Severity

ERROR

### Description

This message occurs when cluster LIF route update fails upon node switchover.

### Corrective Action

Verify that node has connectivity to Amazon Web Service(AWS) Elastic Cloud Computing(EC2) endpoint; check proxy and security group settings. Verify that node has permissions to access AWS EC2 endpoint, check Identity and Access Management (IAM) roles. Verify that any resource is properly specified; check any resource exists and is allowed to be used for specified purpose. Verify that resources have not reached capacity limits.

### Syslog Message

Action: "%s" has failed. HTTP\_Response: "%s"

## Parameters

**action** (STRING): Amazon Web Services (AWS) API command for route table updates or description.  
**error** (STRING): HTTP error response.

## **vsa.cloud.diskReadsThrottled**

### **Severity**

NOTICE

### **Description**

This message occurs when Cloud Volumes ONTAP® detects disk I/O (Reads) throttling on its attached disks. Disk I/O throttling might be initiated by the underlying host or cloud provider when it detects that disk I/Os have reached a certain threshold level mandated per virtual machine instance.

### **Corrective Action**

(None).

### **Syslog Message**

Number of disks detected with read I/O throttling: %s. Interval: (%s) to (%s).

### **Parameters**

**start\_interval** (STRING): Start time of the interval during which disk I/Os were throttled.

**end\_interval** (STRING): End time of the interval during which disk I/Os were throttled.

**disk\_list** (STRING): List of disk device names that were detected with throttled disk I/Os (reads).

## **vsa.cloud.diskWritesThrottled**

### **Severity**

NOTICE

### **Description**

This message occurs when Cloud Volumes ONTAP® detects disk I/O (writes) throttling on its attached disks. Disk I/O throttling might be initiated by the underlying host or cloud provider when it detects that disk I/Os have reached a certain threshold level mandated per virtual machine instance.

### **Corrective Action**

(None).

### **Syslog Message**

Number of disks detected with write I/O throttling: %s. Interval: (%s) to (%s).

### **Parameters**

**disk\_list** (STRING): List of disk device names that detected with throttled disk I/Os (writes).

**start\_interval** (STRING): Start time of the interval during which disk I/Os were throttled.

**end\_interval** (STRING): End time of the interval during which disk I/Os were throttled.

## **vsa.cloud.nlb.probeActive**

### **Severity**

NOTICE

### **Description**

This message occurs when Cloud Volumes ONTAP® receives a probe from a Cloud Network Load Balancer, indicating that a transition to the expected normal state has occurred. The port has gone to active state.

## Corrective Action

(None).

## Syslog Message

Received Cloud Load Balancer probe (now active) for %d ports (port range: %d to %d).

## Parameters

**activated\_ports** (INT): Number of ports that received the probe.

**first\_activated\_port** (INT): First (lowest) port that received a probe.

**last\_activated\_port** (INT): Last (highest) port that received a probe.

## vsa.cloud.nlb.probeInactive

### Severity

ALERT

### Description

This message occurs when Cloud Volumes ONTAP® does not receive expected health probes from the Cloud Network Load Balancer service within the expected time limit. This might cause the Cloud Network Load Balancer to fail over traffic from the corresponding address. If the Cloud Network Load Balancer does not have an active address, all traffic might be blocked from reaching the configured address.

### Corrective Action

Verify the configuration of the Cloud Network Load Balancer, using the cloud provider's console. Verify the ONTAP status of the Cloud Network Load Balancer, as well as the configuration and status of the LIF. To show the ONTAP status for all Cloud Network Load Balancer probe ports, use the (privilege: diagnostic) "virtual-machine show-network-load-balancer" command. To show the ONTAP status for just the inactive Cloud Network Load Balancer probe ports, use the (privilege: diagnostic) "virtual-machine show-network-load-balancer -active false" command. In the above command check if the "Last Probe Time" is being regularly updated (approximately every 5 seconds) and the Cloud Load Balancer probe port matches the "Probe Port".

### Syslog Message

Failed to receive Cloud Load Balancer probe (now inactive) for %d ports (port range: %d to %d), within %d seconds.

### Parameters

**missed\_ports** (INT): Number of ports that did not receive a probe.

**first\_missed\_port** (INT): First (lowest) port that did not receive a probe.

**last\_missed\_port** (INT): Last (highest) port that did not receive a probe.

**timeout\_in\_seconds** (LONGINT): Seconds before timeout occurs.

## vsa.cloud.nlb.probeListening

### Severity

NOTICE

### Description

This message occurs when a Cloud Volumes ONTAP® probe port starts or stops listening for probes from a Cloud Network Load Balancer. The most common scenarios are when node boots up or shuts down.



## Corrective Action

(None).

## Syslog Message

ONTAP Load Balancer has %s listening for probes from a Cloud Network load Balancer on port %d.

## Parameters

**detail** (STRING): Whether the Cloud Volumes ONTAP probe port has started or stopped listening for probes from the Cloud Load Balancer.

**probe\_port** (INT): Port that the listening event occurred on.

# vsa.device events

## vsa.device.bad.config

### Severity

ALERT

### Description

This message occurs when the system detects that the kernel environment variables (kenvs) that are used to configure the core dump and NVRAM devices are missing or inconsistent. If the system is rebooted with a bad configuration, this may result in a failure to boot or incorrect operation of NVRAM or core dump.

### Corrective Action

Correct the configuration as directed. Use the "kenv" systemshell command to set and verify the environment variables. Then reboot the system. If the node is not booting or is in a boot loop, use the Loader commands to set and verify the kernel environment variables before rebooting.

### Syslog Message

Core dump or NVRAM device configuration problem: %s.

### Parameters

**detail** (STRING): Detail about the misconfiguration and suggestions for correcting the problem.

# vsa.diskobjectstore events

## vsa.diskObjectStore.cfg.sync

### Severity

ALERT

### Description

This message occurs when the system detects that a disk object store configuration is not synchronized between both nodes in a shared-storage High-Availability (HA) pair.

### Corrective Action

Use the "system node virtual-machine disk-object-store" commands to synchronize the object store configuration.

## Syslog Message

Object store configuration %s synchronization issue. Object store "%s" on node "%s" is not synchronized.  
Details: %s.

## Parameters

**detail** (STRING): The disk object store configuration for shared disks must be identical on both nodes in an HA pair. The system attempts to keep them synchronized when the object store configuration is created, deleted, or modified. However, under unusual circumstances, the synchronization may be unsuccessful. Manual intervention is required to synchronize the configuration.

**object\_store\_name** (STRING): Object store name.

**partner** (STRING): Partner node name.

**synchronization\_details** (STRING): Synchronization state details. For example, the configuration that is missing on the partner, or a list of fields that do not have the same value.

# vsa.ephemeral events

## vsa.ephemeral.vnvram.reset

### Severity

NOTICE

### Description

This message occurs when a node is configured to use ephemeral vNVRAM and the virtual machine has been re-hosted, using new ephemeral disk for vNVRAM.

### Corrective Action

(None).

## Syslog Message

Ephemeral vNVRAM disk for node has changed, due to re-hosting of VM. New serial number: "%s".

## Parameters

**ephemeral\_vnvram\_serial\_number** (STRING): Value of the serial number of the updated ephemeral disk used for vNVRAM.

# vsa.instancemetadata events

## vsa.instanceMetadata.change

### Severity

ALERT

### Description

This message occurs when the system detects that a field in the virtual machine's instance metadata settings has changed. The virtual machine instance metadata holds configuration settings for the instance. Some virtual machine settings can be changed after deployment, and can be changed using cloud provider tools. These settings include some that can affect the performance or stability of the instance. One example is the instance type. If the instance type is changed to an unsupported value, this might result in hard-to-diagnose stability issues if that instance type does not have enough memory or CPU cores. This message is generated only for changes that might impact the performance or reliability of the instance.

### Corrective Action

If the field is changed to an unsupported value, manual reconfiguration will be required.

### Syslog Message

Virtual machine instance metadata setting %s. Field: %s. Previous value: "%s". New value: "%s".

### Parameters

**detail** (STRING): Detail about how the change affects the instance.

**field\_name** (STRING): Instance metadata field name.

**previous\_value** (STRING): Previous value of field.

**current\_value** (STRING): Current value of field.

## vsa.mlx events

### vsa.mlx.nic.detach

#### Severity

NOTICE

#### Description

This message occurs when ONTAP® software is running as an instance in Azure and a NIC SR-IOV Mellanox (MLX) driver detach event is detected. ONTAP will continue to use the Hyper-V netsh driver without the SR-IOV functionality for a short time. On HA configurations, unsynchronized logs might appear during this time, which will be resolved if a subsequent SR-IOV attach event occurs. If the SR-IOV attach event does not occur quickly, high-availability (HA) takeover might occur.

#### Corrective Action

(None).

#### Syslog Message

NIC SR-IOV driver detach event detected. On HA configurations, unsynchronized logs might result or an HA takeover might occur if SR-IOV remains detached.

#### Parameters

(None).

## vsa.nodesernum events

### vsa.nodeSerNum.change.failed

#### Severity

ALERT

#### Description

This message occurs when a node serial number change operation fails. The serial number might need to be reverted to the previous value, and the node might need to be rebooted.

#### Corrective Action

Check the node serial number and state of the node licensing. The serial number might need to be reverted to the previous value. If the reversion does not work, or if there are licensing issues reported by the system,

reboot the node. After the node is online, check the serial number and licensing for any issues.

### **Syslog Message**

Node serial number change operation failed. Error detail: %s. Previous value: "%s". New value: "%s".

### **Parameters**

**error\_detail** (STRING): Error details, if available.

**previous\_value** (STRING): Previous value of the serial number.

**new\_value** (STRING): New value of the serial number.

## **vsa.nodeSerNum.changed**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a node serial number change operation is completed successfully.

### **Corrective Action**

(None).

### **Syslog Message**

Node serial number change operation was completed successfully. Previous value: "%s". New value: "%s".

### **Parameters**

**previous\_value** (STRING): Previous value of the serial number.

**new\_value** (STRING): New value of the serial number.

## **vsa.savecore events**

### **vsa.savecore.nospace**

### **Severity**

ERROR

### **Description**

This message occurs when the savecore tool determines that the space available is not sufficient for storing the core image from a dump device to a readable core file in the /etc/crash directory.

### **Corrective Action**

Try performing the following steps to make sufficient room in root volume: Delete old or unneeded core files by using the "coredump delete" command. Delete old or unneeded core file segments by using the "coredump segment delete" command. Delete any Snapshot® copies that include the deleted files. Adjust minimum free space on the root file system by using "coredump config modify" command. Once there is sufficient room, take the following steps to save the cores: To save all core files using "coredump save-all" command, run the "coredump status -instance" command and ensure that "Space Available On Internal Filesystem" is greater than "Space Needed to Save All Unsaved Cores" plus "Minimum Free Bytes on Root Filesystem". To save a specific core file, run the "coredump save" command, when available space of "Space Needed To Save Core" (from the "coredump show" command) plus "Minimum Free Bytes on Root Filesystem" is sufficient.

## Syslog Message

Available space (%llu %dKB blocks) is not enough to save the core image (requires up to %llu blocks) and maintain the minimum free space of %llu blocks.

## Parameters

**avail** (LONGINT): Number of blocks available for saving the core image.

**blocksize** (INT): Block size in kilobytes.

**needed** (LONGINT): Number of blocks needed for saving the core image.

**minfree** (LONGINT): Number of blocks reserved for normal system operation.

# vsa.scheduledevent events

## vsa.scheduledEvent.mlx.detach

### Deprecated

This event is replaced with `vsa.mlx.nic.detach` which is the same, but with a new non-confusing name.

### Severity

NOTICE

### Description

This message occurs when ONTAP® software is running as an instance in Azure and a NIC SR-IOV driver detach event is detected. ONTAP will continue to use the Hyper-V netsh driver without the SR-IOV functionality for a short time. On an HA configuration, unsynchronized logs might appear during this time, which will be resolved if a subsequent SR-IOV attach event occurs. If the SR-IOV attach event does not occur quickly, high-availability (HA) takeover might occur.

### Corrective Action

(None).

### Syslog Message

NIC SR-IOV driver detach event detected. On an HA configuration, unsynchronized logs might be seen or HA takeover might occur if SR-IOV driver remains detached.

### Parameters

(None).

## vsa.scheduledEvent.scheduled

### Severity

ALERT

### Description

This message occurs when the system detects that a cloud provider event has been scheduled that will result in a node rebooting or becoming unavailable. In a High-Availability (HA) configuration, only one partner will be affected.

### Corrective Action

In a single-node configuration, a data-serving failure will occur for an event of type "reboot" or "rehost", and a data-serving failure or timeout might occur during a "freeze". Prepare client applications for the event. For an HA configuration, no action is required.

## Syslog Message

Cloud provider event %s for node "%s". ID: %s. Type: %s. Not before: %s.

### Parameters

**detail** (STRING): Cloud provider infrastructure might cause nodes to reboot or freeze. This happens when the provider needs to update the underlying infrastructure or move the node to a new host. Typically, the event will be detected before it happens, and information about when it might start is provided via a "not before" time. In an HA configuration, only one node in an HA pair will be affected by an event. During this time, the cluster will not be protected from a data-serving failure if the remaining HA partner goes down.

**node** (STRING): Node affected by the event.

**event\_id** (STRING): Cloud-provider-specific event ID.

**event\_type** (STRING): String "reboot", "rehost", or "freeze".

**not\_before\_time** (STRING): Timestamp in the general form: Mon, 19 Sep 2018 18:29:47 GMT If this string is empty, the event is in progress.

## vsa.scheduledEvent.update

### Severity

NOTICE

### Description

This message occurs when the system detects that the status of a previously scheduled cloud provider maintenance event has been updated.

### Corrective Action

Cloud provider infrastructure will restart the node after a "reboot" or "rehost" event, and during a "freeze" the node might restart. In any case, intervention might be required to be sure that the cluster is up and running normally.

## Syslog Message

Cloud provider event %s for node "%s". ID: %s. Type: %s. Status: %s.

### Parameters

**detail** (STRING): Cloud provider maintenance events have a lifecycle. When the event is scheduled, the "scheduled" event is raised. Subsequently, the status of the event will change state to "started" or "complete". When that is detected, this event is raised.

**node** (STRING): Node affected by the event.

**event\_id** (STRING): Cloud-provider-specific event ID.

**event\_type** (STRING): String "reboot", "rehost", or "freeze".

**status** (STRING): Current status of the event: "scheduled", "started", "complete".

## vsa.vsphere events

### vsa.vsphere.connected

#### Severity

INFORMATIONAL

#### Description

This message occurs when Data ONTAP® on VMware® manages to reconnect and authenticate itself with the vSphere(tm) server.

**Corrective Action**

None.

**Syslog Message**

Authenticated connection (user %s) to vSphere server %s reestablished.

**Parameters**

**user** (STRING): vSphere "read-only" user name.

**server** (STRING): Host name of the vSphere server.

**vsa.vsphere.noauth****Severity**

ERROR

**Description**

This message occurs when Data ONTAP® on VMware® cannot authenticate itself to the vSphere(tm) server to gather host hardware information.

**Corrective Action**

Correct the vSphere server authentication information with the command 'vsphere modify'.

**Syslog Message**

Cannot authenticate with vSphere server: %s.

**Parameters**

**detail** (STRING): Details of how authentication failed.

**vsa.vsphere.unreachable****Severity**

ERROR

**Description**

This message occurs when Data ONTAP® on VMware® cannot connect to the vSphere(tm) server to gather host hardware information.

**Corrective Action**

Verify that the vSphere server is up and reachable through https: If necessary, correct vSphere server information using the command 'vsphere modify'.

**Syslog Message**

Cannot establish connection to vSphere server: %s.

**Parameters**

**detail** (STRING): Details of how the connection failed.

# vscan events

## vscan.disabled events

### vscan.disabled

**Severity**

INFORMATIONAL

**Description**

This message occurs when a Vserver or cluster administrator disables Vscan on a Vserver. No subsequent client requests will trigger virus scanning.

**Corrective Action**

(None).

**Syslog Message**

Vscan is disabled on Vserver '%s'.

**Parameters**

**vserverName** (STRING): Name of the Vserver.

## vscan.enabled events

### vscan.enabled

**Severity**

INFORMATIONAL

**Description**

This message occurs when a Vserver or cluster administrator enables Vscan on a Vserver. Based on the configuration, subsequent client requests can trigger virus scanning.

**Corrective Action**

(None).

**Syslog Message**

Vscan is enabled on Vserver '%s'.

**Parameters**

**vserverName** (STRING): Name of the Vserver.

## vscan.newversion events

### vscan.newVersion.allocated

**Severity**

INFORMATIONAL



## Description

This message occurs when the Vscan version mechanism allocates a new version-ID corresponding to the Vscan server version.

## Corrective Action

(None).

## Syslog Message

Vscan version mechanism added new version-ID for Vserver '%s' corresponding to vendor '%s', version '%s'.

## Parameters

**vservername** (STRING): Name of the Vserver.

**vendor** (STRING): Vscan server vendor name.

**version** (STRING): Vscan server version.

# vscan.odjob events

## vscan.odJob.delete.failed

### Severity

ERROR

### Description

This message occurs when the system fails to delete queued Vscan On-Demand job. The system tries to delete this job internally in a situation such as when Vserver is stopped or when a Vserver is failed over to some other cluster.

### Corrective Action

Run the "vserver vscan on-demand-task unschedule" command to delete the job.

### Syslog Message

System failed to delete Vscan On-Demand job for the task "%s" on Vserver "%s". Reason: "%s".

### Parameters

**taskName** (STRING): Vscan On-Demand task name.

**vservername** (STRING): Name of the Vserver.

**failureReason** (STRING): Reason for the failure.

## vscan.odJob.running.deleted

### Severity

NOTICE

### Description

This message occurs when the system deletes a running On-Demand job internally.

### Corrective Action

(None).

## Syslog Message

System deleted a running Vscan On-Demand job. Job ID "%s".

## Parameters

**jobID** (STRING): Job ID.

## vscan.odJob.schedule.failed

### Severity

ERROR

### Description

This message occurs when the system fails to schedule a previously requested Vscan On-Demand job. The system tries to schedule this job internally in a situation such as when a Vserver is started or when a Vserver is failed over to some other cluster.

### Corrective Action

Run the "vserver vscan on-demand-task schedule" command to schedule the job.

## Syslog Message

System failed to schedule Vscan On-Demand job for the task "%s" on Vserver "%s". Reason: "%s".

## Parameters

**taskName** (STRING): Vscan On-Demand task name.

**vserverName** (STRING): Name of the Vserver.

**failureReason** (STRING): Reason for the failure.

## vscan.odreport events

### vscan.odReport.create.failed

#### Severity

ERROR

#### Description

This message occurs when the On-Demand job fails to create report file.

#### Corrective Action

Ensure that the report directory exists and is accessible.

#### Syslog Message

On-Demand job for the task "%s" on Vserver "%s" could not create report file in the directory "%s". Reason: "%s".

#### Parameters

**taskName** (STRING): Vscan On-Demand task name.

**vserverName** (STRING): Name of the Vserver.

**reportDirectory** (STRING): Report directory.

**failureReason** (STRING): Reason for the failure.

## **vscan.odReport.write.failed**

### **Severity**

ERROR

### **Description**

This message occurs when the On-Demand job fails to write to a report file.

### **Corrective Action**

Ensure that the report file is accessible and the volume on which it resides has space.

### **Syslog Message**

On-Demand job for the task "%s" on Vserver "%s" could not write to the report file "%s". Reason: "%s".

### **Parameters**

**taskName** (STRING): Vscan On-Demand task name.

**vserverName** (STRING): Name of the Vserver.

**reportFile** (STRING): Path of the report file.

**failureReason** (STRING): Reason for the failure.

## **vscan.pool events**

### **vscan.pool.autoActivated**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the Vscan scanner pool is automatically activated.

#### **Corrective Action**

(None).

#### **Syslog Message**

Vscan scanner pool '%s' is automatically activated on Vserver '%s'.

#### **Parameters**

**scannerPool** (STRING): Name of the scanner pool.

**vserverName** (STRING): Name of the Vserver.

### **vscan.pool.autoDeactivated**

#### **Severity**

INFORMATIONAL

#### **Description**

This message occurs when the Vscan scanner pool is automatically deactivated.

#### **Corrective Action**

(None).

## Syslog Message

Vscan scanner pool '%s' is automatically deactivated on Vserver '%s'.

## Parameters

**scannerPool** (STRING): Name of the scanner pool.

**vserverName** (STRING): Name of the Vserver.

## vscan.pool.largeReqTimeout

### Severity

NOTICE

### Description

This message occurs when the request timeout of a Vscan scanner pool is configured more than what can be honored for an On-Access policy whose scan-mandatory option is set to "off". The system internally adjusts the request timeout to the maximum allowed value.

### Corrective Action

(None).

## Syslog Message

For scanner-pool '%s' created on Vserver '%s', request-timeout is configured as %ds but that would internally be taken as %ds if the applicable On-Access policy has scan-mandatory set to "off". This ensures that the file-access is granted to the client for such policy.

## Parameters

**scannerPool** (STRING): Name of the scanner pool.

**vserverName** (STRING): Name of the Vserver on which the scanner pool is configured.

**requestTimeout** (INT): Request timeout value configured by the user.

**adjustedRequestTimeout** (INT): Request timeout value that the system internally adjusts to.

## vscan.privsharecreate events

### vscan.privShareCreate.failed

### Severity

ERROR

### Description

This message occurs when the creation of privileged share ONTAP\_ADMIN\$ fails. Attempted connections to ONTAP\_ADMIN\$ share by the Vscan servers will fail.

### Corrective Action

Check the reason for failure and take action accordingly. If the problem persists, contact NetApp technical support for assistance.

## Syslog Message

Failed to create privileged share ONTAP\_ADMIN\$ for Vserver "%s". Reason: "%s".

## Parameters

**vserverName** (STRING): Name of the Vserver.

**reason** (STRING): Reason for failure.

# vscan.rdbupdregister events

## vscan.rdbUpdRegister.failed

### Severity

ERROR

### Description

This message occurs when the Vscan configuration replication mechanism fails to register RDB update callback. Modifications made to the Vscan configuration might not be available to this node.

### Corrective Action

Perform a takeover and giveback of the node or reboot the node. For additional options, contact NetApp technical support.

### Syslog Message

Vscan configuration replication mechanism failed to register RDB update callback. Modifications made to the Vscan configuration might not be available to this node.

## Parameters

(None).

# vsdr events

## vsdr.adtcreate events

### vsdr.adtcreate.fail

#### Severity

ERROR

#### Description

This message occurs when an audit staging volume cannot be created at the destination cluster in Async-DR (asynchronous disaster recovery), either because there is not enough disk space on the aggregates or the aggregates are in offline state.

#### Corrective Action

If the aggregates are offline, bring them online. Ensure that each aggregate has at least 2GB of free space, adding disks, reducing volume size or moving volumes as needed, and then try the following commands: Then try the following commands: `prompt> snapmirror initialize` `prompt> snapmirror resync`

#### Syslog Message

Vserver audit create failed at the destination cluster with Async-DR.

#### Parameters

(None).

## vsdr.app events

### vsdr.app.cg.conflict

#### Severity

NOTICE

#### Description

This message occurs during the configuration replication phase of a Vserver DR transfer when the consistency group that the system is trying to replicate already exists in the cluster on a different Vserver.

#### Corrective Action

Remove the existing consistency group or make the Vserver that has the conflicting consistency group in it the destination in the Vserver DR relationship.

#### Syslog Message

Consistency group %s cannot be created in Vserver %s due to conflicting consistency group %s in Vserver %s.

#### Parameters

**consistency\_group** (STRING): Name of the consistency group that the system is trying to replicate.

**destination\_vserver** (STRING): Name of the destination Vserver into which the consistency group is being replicated.

**conflicting\_consistency\_group** (STRING): Name of the conflicting consistency group.

**found\_in\_vserver** (STRING): Name of the Vserver in which the conflicting consistency group was found.

# vsdr.clone events

## vsdr.clone.cleanup.fail

### Severity

ERROR

### Description

This message occurs when a clone create job fails prematurely. As result, there is a partially created clone in the cluster.

### Corrective Action

Use the (privilege: diag) "debug vreport show -type volume" command to show the list of inconsistent volumes. For each of these inconsistent volumes, use the (privilege: diag) "debug vreport fix -type volume" command to resolve all of the issues. If this is a SnapMirror® or MetroCluster(tm) environment, run the "snapmirror resync" command or the "metrocluster vservers resync" command, respectively, after the inconsistencies are resolved.

### Syslog Message

Clone create job failed prematurely for volume "%s" (Vserver Id %d). Reason:%s.

### Parameters

**volume** (STRING): Name of the clone volume for which the create operation failed.

**vserver\_id** (INT): ID of the Vserver that owns the volume.

**errmsg** (STRING): Error message string.

## vsdr.clone.splitOnSource

### Severity

ERROR

### Description

This message occurs when a clone volume is split on the source Vserver in a Vserver DR setup, and a SnapMirror® transfer is initiated after this operation.

### Corrective Action

Run the "volume clone split start" command to split the volume on the dp-destination Vserver. Then run the "volume delete" command to remove the volume from the dp-destination Vserver. Then run the "snapmirror resync" command on the VserverDR relationship.

### Syslog Message

The volume "%s" of Vserver "%s" has been split on the source Vserver. Split the clone volume on the destination as well. After splitting, delete the volume and execute a "snapmirror resync" command.

### Parameters

**volume\_name** (STRING): Volume clone that needs to be split.

**vserver\_name** (STRING): Vserver name for which the SnapMirror transfer has been initiated.

# vsdr.clonetmp events

## vsdr.clonetmp.msids.conflict

### Severity

ERROR

### Description

This message occurs during a SnapMirror® resync operation on a Storage Virtual Machine (SVM) disaster recovery (DR) relationship, when the MSID of a volume on the source SVM is the same as the MSID of a temporary volume on the destination SVM and the temporary volume is created for a clone parent volume. The resync operation fails to start.

### Corrective Action

Split the clone volume by using the "volume clone split start" command or delete the clone volume by using the "volume delete" command, and then remove the temporary volume by using the "volume delete" command before executing another "snapmirror resync" command.

### Syslog Message

The MSID of temporary clone parent volume "%s" on the destination SVM "%s" conflicts with the MSID of volume "%s" on the source".

### Parameters

**tmp\_volume\_name** (STRING): Volume name of the temporary volume on the destination SVM that is a parent of a clone volume.

**vserver\_name** (STRING): SVM name for which the SnapMirror resync was initiated.

**src\_volume\_name** (STRING): Volume name on the source SVM.

# vsdr.cluster events

## vsdr.cluster.sourceVersionIsHigherThanDestination

### Severity

ERROR

### Description

This message occurs when a source cluster of SVMDR relationship is at higher version than destination cluster.

### Corrective Action

Upgrade source cluster and destination cluster to the same version.

### Syslog Message

Resync failed because source cluster version "%s" is higher than destination cluster version "%s"

### Parameters

**source\_cluster\_version** (STRING): Version of source cluster of SVMDR relationship.

**destination\_cluster\_version** (STRING): Version of destination cluster of SVMDR relationship.



# vsdr.fpvolallpolicy events

## vsdr.FPVolAllPolicy

### Severity

ALERT

### Description

This message occurs when SnapMirror® transfer happens between the source volume of an SVM disaster recovery (SVMDR) relationship having a tiering policy of "all" residing on a FabricPool aggregate on the primary cluster and destination volume residing on a FabricPool aggregate at the secondary cluster. SnapMirror transfer will be slow and might cause recovery point objective (RPO) to be missed.

### Corrective Action

Move the source volume of the SVMDR relationship from the FabricPool aggregate to a non-FabricPool aggregate at the primary cluster or modify the volume tiering policy of the source volume. To move the source volume, use the "volume move start -vserver <vserver-name> -volume <volume-name> -destination -aggregate <aggregate> command." To modify the volume tiering policy, use the "volume modify -volume <volume-name> -vserver <vserver-name> -tiering-policy <snapshot-only|none|auto>" command.

### Syslog Message

SnapMirror® transfer between source volume "%s" on Vserver "%s" on a FabricPool aggregate having a tiering policy of "all" on the primary cluster and the destination volume on a FabricPool aggregate at the destination cluster will be slow.

### Parameters

**source\_volume\_name** (STRING): Name of the source volume on FabricPool aggregate of SVMDR relationship.

**src\_vserver\_name** (STRING): Name of the source vserver name.

# vsdr.fpvolallpolicyonfpaggr events

## vsdr.FPVolAllPolicyOnFPAggr

### Severity

ALERT

### Description

This message occurs when source volume of an SVM disaster recovery (SVMDR) relationship having a volume tiering policy of "all" residing on FabricPool aggregate on the primary cluster is placed on a FabricPool aggregate at the destination cluster. SnapMirror® transfer will be slow and might cause recovery point objective(RPO) to be missed.

### Corrective Action

Either move the source volume of the SVMDR relationship from the FabricPool aggregate to a non-FabricPool aggregate at the primary cluster or modify the volume tiering policy of the source volume. To move the source volume, use the "volume move start -vserver <vserver-name> -volume <volume-name> -destination-aggregate <aggregate> command." To modify the volume tiering policy, use the "volume modify -volume <volume-name> -vserver <vserver-name> -tiering-policy <snapshot-only|none|auto>" command.

## Syslog Message

Source volume "%s" of the SVMDR relationship residing on a FabricPool aggregate having a tiering policy of "all" on the primary cluster is placed on FabricPool aggregate "%s" at the destination cluster.

## Parameters

**source\_volume\_name** (STRING): Name of the source volume on the FabricPool aggregate of the SVMDR relationship.

**dest\_aggr\_name** (STRING): Name of the destination FabricPool aggregate at the secondary cluster.

# vsdr.fpvoldropqospolicy events

## vsdr.FPVoldDropQosPolicy

### Severity

NOTICE

### Description

This message occurs when the system clears or drops an incompatible quality of service (QoS) policy group setting with minimum throughput value set on a destination volume having tiering policy other than "none" on a FabricPool aggregate at the destination cluster.

### Corrective Action

(None).

## Syslog Message

QoS policy group with minimum throughput setting cannot be set on the volume "%s" on storage virtual machine (SVM) "%s" on a FabricPool aggregate having a tiering policy other than "none".

## Parameters

**source\_volume\_name** (STRING): Name of the source volume on FabricPool aggregate of SVM DR relationship.

**src\_vserver\_name** (STRING): Name of the source SVM.

# vsdr.fpvolonnonfpaggr events

## vsdr.FPVolOnNonFPAggr

### Severity

NOTICE

### Description

This message occurs when a source volume of an SVM disaster recovery (SVMDR) relationship residing on a FabricPool aggregate at the primary cluster is placed on a non-FabricPool aggregate at the secondary cluster.

### Corrective Action

Move the source volume of SVMDR relationship from the FabricPool aggregate to a non-FabricPool aggregate at the primary cluster by using the "volume move start -vserver <vserver-name> -volume <volume-name> -destination-aggregate <aggregate>" command.

## Syslog Message

Source volume "%s" of the SVMDR relationship residing on a FabricPool aggregate on the primary cluster is placed on a non-FabricPool aggregate "%s" at the destination cluster.

## Parameters

**source\_volume\_name** (STRING): Name of the source volume on the FabricPool aggregate of the SVMDR relationship.

**dest\_aggr\_name** (STRING): Name of the destination non-FabricPool aggregate at the secondary cluster.

# vsdr.invalid events

## vsdr.invalid.license

### Severity

ERROR

### Description

This message occurs when a package is licensed at the primary site but not licensed at the secondary site in a MetroCluster(tm) configuration or Vserver DR setup.

### Corrective Action

Install at the secondary site a license that is valid for the package. If this is a SnapMirror® or MetroCluster(tm) environment, run the "snapmirror resync" command or the "metrocluster vserver resync" command respectively to resynchronize the Vserver configuration.

## Syslog Message

Configuration replication on vserver %s failed due to an invalid license for package %s.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**package\_name** (STRING): Name of the package for which the license is to be added.

# vsdr.ip events

## vsdr.ip.lif.auto.created

### Severity

NOTICE

### Description

This message occurs when a logical interface (LIF) is automatically created through configuration replication in a MetroCluster™ configuration.

### Corrective Action

(None).

## Syslog Message

LIF %s in Vserver %s was created on port %s on node %s.

## Parameters

**lif** (STRING): LIF name.  
**vserver\_name** (STRING): Vserver name.  
**port\_name** (STRING): Name of the home port of the LIF.  
**node\_name** (STRING): Name of the home node of the LIF.

## vsdr.ip.lif.home.changed

### Severity

NOTICE

### Description

This message occurs when the home port or node of an existing IP logical interface (LIF) is modified during LIF placement in a MetroCluster™ configuration.

### Corrective Action

(None).

### Syslog Message

Home port and node changed from %s, %s to %s, %s while placing LIF %s in Vserver %s.

## Parameters

**old\_port\_name** (STRING): Name of the previous home port of the existing IP LIF.  
**old\_node\_name** (STRING): Name of the previous home node of the existing IP LIF.  
**port\_name** (STRING): Name of the current home port of the existing IP LIF.  
**node\_name** (STRING): Name of the current home node of the existing IP LIF.  
**lif** (STRING): LIF name.  
**vserver\_name** (STRING): Vserver name.

## vsdr.missing events

### vsdr.missing.clone.parent.volume

### Severity

ERROR

### Description

This message occurs when the clone parent volume does not exist during the Vserver SnapMirror®/MetroCluster(tm) baseline replication on the secondary site. Vserver SnapMirror/MetroCluster replication will be blocked until this problem is resolved, and delaying action on this problem will lead to inconsistent data on the primary and secondary sites.

### Corrective Action

Use the (privilege: diag) "debug vreport show -type volume" command on the primary site to show the list of inconsistent volumes. For each of these inconsistent volumes, use the (privilege: diag) "debug vreport fix -type volume" command to resolve all of the issues on the Vserver SnapMirror/MetroCluster primary site. After the above problem is resolved, run the "snapmirror resync/metrocluster vserver resync" command.

### Syslog Message

Parent volume for clone volume %s:%s is missing.

## Parameters

**vserver** (STRING): Name of the Vserver.

**volume** (STRING): Name of the clone volume for which the corresponding parent volume is missing.

# vsdr.net events

## vsdr.net.IPInDiffVserver

### Severity

ERROR

### Description

This message occurs when the system tries to automatically replicate a logical interface (LIF) in a MetroCluster™ configuration or Vserver DR setup and the IP address of the LIF that is being created already exists in another Vserver.

### Corrective Action

Change the IP address of the LIF in the other Vserver or remove the LIF so that creation of the LIF succeeds and configuration replication continues.

### Syslog Message

Cannot create LIF %s with IP address %s in Vserver %s because another LIF with the same IP exists in another Vserver.

## Parameters

**lif** (STRING): LIF name.

**ipaddr** (STRING): IP address of the LIF.

**vserver\_name** (STRING): Vserver name.

## vsdr.net.LifIpCollision

### Severity

ERROR

### Description

This message occurs in an ID-Preserve VserverDR configuration of clustered Data ONTAP®, when a non-SAN LIF created on the source Storage Virtual Machine(SVM) has the same IP address as an iSCSI LIF on the destination SVM.

### Corrective Action

Assign a unique IP address for the iSCSI LIF on the destination SVM.

### Syslog Message

An iSCSI LIF is configured in the destination SVM %s with the same IP address (%s) as the source LIF (%s). Assign a unique IP address to the iSCSI LIF (%s) by modifying the LIF on the destination SVM.

## Parameters

**vserver\_name** (STRING): SVM name.

**ip\_address** (STRING): IP address of the iSCSI LIF on the destination SVM.

**source\_lif\_name** (STRING): Source LIF name.

**destination\_lif\_name** (STRING): Destination LIF name.

## **vsdr.net.LifNameCollision**

### **Severity**

ERROR

### **Description**

This message occurs in an ID-Preserve VserverDR configuration of clustered Data ONTAP®, when a non-SAN LIF created on the source Storage Virtual Machine(SVM) has the same name as a SAN LIF on the destination SVM.

### **Corrective Action**

Assign a unique LIF name for the SAN LIF on the destination SVM.

### **Syslog Message**

A SAN LIF is configured in the destination SVM %s with the same name as the source LIF. Assign a unique name to the SAN LIF (%s) by modifying the LIF name on the destination SVM.

### **Parameters**

**vserver\_name** (STRING): SVM name.

**lif\_name** (STRING): Source and Destination LIF name.

## **vsdr.net.LifNotReachable**

### **Severity**

ERROR

### **Description**

This message occurs when a logical interface (LIF) modify operation is performed on the source cluster without changing the home node and home port, and the LIF placement fails when the configuration is being replicated.

### **Corrective Action**

Ensure that the destination cluster has ports that have connectivity to the LIF on the source cluster.

### **Syslog Message**

The LIF %s in DR partner Vserver of Vserver %s is not reachable from the local cluster.

### **Parameters**

**lif** (STRING): The LIF name.

**vserver\_name** (STRING): The Vserver name.

## **vsdr.net.NoPortsInBD**

### **Severity**

ERROR

### **Description**

This message occurs when the system tries to automatically replicate a logical interface (LIF) in a MetroCluster™ configuration or Vserver DR setup, and it cannot find any ports in a broadcast domain in a cluster.

## Corrective Action

Add at least one port to the broadcast domain in the IPspace using the "network port broadcast-domain add-ports" command and then resynchronize the Vserver configuration using the "metrocluster vserver resync" command.

## Syslog Message

Cannot find ports in broadcast domain %s in IPspace %s for placing LIF %s in Vserver %s.

## Parameters

**broadcast\_domain** (STRING): The Broadcast domain.

**ipspace\_name** (STRING): The IPspace name.

**lif** (STRING): The LIF name.

**vserver\_name** (STRING): The Vserver name.

## vsdr.net.NoPortsInIpsOnNode

### Severity

ERROR

### Description

This message occurs when the system tries to automatically replicate a logical interface (LIF) in a MetroCluster™ configuration or Vserver DR setup, and it cannot find any ports in an IPspace on a particular node.

### Corrective Action

Add at least one port belonging to the node to the IPspace by using the "network port broadcast-domain add-ports" command and then resynchronize the Vserver configuration using the "metrocluster vserver resync" command.

### Syslog Message

Cannot find ports in IPspace %s on node %s for placing LIF %s in Vserver %s.

### Parameters

**ipspace\_name** (STRING): The IPspace name.

**node\_name** (STRING): Name of the node to which LIFs could not be replicated.

**lif** (STRING): The LIF name.

**vserver\_name** (STRING): Vserver name.

## vsdr.net.NoPortsInIpspace

### Severity

ERROR

### Description

This message occurs when the system tries to automatically replicate a logical interface (LIF) in a MetroCluster™ configuration or Vserver DR setup, and it cannot find any ports in an IPspace in a cluster.

### Corrective Action

Add at least one port to the IPspace by using the "network port broadcast-domain add-ports" command and then resynchronize the Vserver configuration using the "metrocluster vserver resync" command.

## Syslog Message

Cannot find ports in IPspace %s for placing LIF %s in Vserver %s.

## Parameters

**ipspace\_name** (STRING): The IPspace name.

**lif** (STRING): The LIF name.

**vserver\_name** (STRING): The Vserver name.

## vsdr.net.NoVipPortInIpsNode

### Severity

ERROR

### Description

This message occurs when the system attempts to replicate a logical interface (LIF) associated with the Virtual IP (VIP) feature to a node where no BGP peer group is available in the target IPspace. This configuration replication functionality is exercised by systems in a MetroCluster(tm).

### Corrective Action

Use the "network bgp peer-group create" command to create a BGP peer group in the target IPspace on the affected node, and then use the "metrocluster vserver resync" command to recover from the failed replication.

## Syslog Message

Cannot find a VIP port in IPspace %s on node %s for placing VIP LIF %s in Vserver %s.

## Parameters

**ipspace\_name** (STRING): IPspace name.

**node\_name** (STRING): Name of the node to which VIP LIFs could not be replicated.

**lif** (STRING): VIP LIF name.

**vserver\_name** (STRING): Vserver name.

## vsdr.net.NoVipPortInIpspace

### Severity

ERROR

### Description

This message occurs when the system attempts to replicate a logical interface (LIF) associated with the Virtual IP (VIP) feature in a MetroCluster™ configuration or Vserver DR setup, and the Border Gateway Protocol (BGP) is not configured in the target IPspace.

### Corrective Action

Use the "network bgp peer-group create" command to create a BGP peer group in the target IPspace, and then use the "metrocluster vserver resync" command to recover from the failed replication.

## Syslog Message

Cannot find a VIP port in IPspace %s for placing VIP LIF %s in Vserver %s.

## Parameters

**ipspace\_name** (STRING): IPspace name.



**lif** (STRING): VIP LIF name.

**vserver\_name** (STRING): Vserver name.

## **vsdr.net.v6optdisabled**

### **Severity**

ERROR

### **Description**

This message occurs in a MetroCluster™ configuration when the system is trying to replicate an IPv6 LIF, but IPv6 is not enabled on the cluster.

### **Corrective Action**

Enable IPv6 by using the "network options ipv6 modify -enabled true" command, and then run the "metrocluster vserver resync" command on the source cluster.

### **Syslog Message**

Cannot create LIF %s in Vserver %s because IPv6 is not enabled in the cluster.

### **Parameters**

**lif** (STRING): LIF name.

**vserver\_name** (STRING): Vserver name.

## **vsdr.net.v6optdisabledroute**

### **Severity**

ERROR

### **Description**

This message occurs when the system is trying to replicate an IPv6 route in a MetroCluster™ configuration, but IPv6 is not enabled on the cluster.

### **Corrective Action**

Enable IPv6 by using the "network options ipv6 modify -enabled true" command, and then run the "metrocluster vserver resync" command on the source cluster.

### **Syslog Message**

Cannot create route %s with gateway %s in Vserver %s because IPv6 is not enabled in the cluster.

### **Parameters**

**route** (STRING): Route destination and mask.

**gateway** (STRING): Gateway of the route.

**vserver\_name** (STRING): Vserver name.

## **vsdr.nonfpvolonfpaggr events**

### **vsdr.nonFPVolOnFPAggr**

### **Severity**

NOTICE

## Description

This message occurs when the source volume of an SVM disaster recovery (SVMDR) relationship residing on a non-FabricPool aggregate at the primary cluster is placed on a FabricPool aggregate at the secondary cluster.

## Corrective Action

Move the source volume of the SVMDR relationship from the non-FabricPool aggregate to a FabricPool aggregate at the primary cluster by using the "volume move start -vserver <vserver-name> -volume <volume-name> -destination-aggregate <aggregate>" command.

## Syslog Message

Source volume "%s" of the SVMDR relationship residing at a nonFabricPool aggregate at the primary cluster is placed on a FabricPool aggregate "%s" at the destination cluster.

## Parameters

**source\_volume\_name** (STRING): Name of the source volume on the non-FabricPool aggregate of the SVMDR relationship.

**dest\_aggr\_name** (STRING): Name of the destination FabricPool aggregate at the secondary cluster.

# vsdr.periodic events

## vsdr.periodic.check.errors

### Severity

ERROR

### Description

This message occurs when the Vserver DR subsystem periodic check detects issues with the Vserver DR components. These issues could potentially result in a degraded Vserver DR relationship. So, some corrective action is required by the admin to resolve the issue.

### Corrective Action

Contact NetApp technical support to identify and resolve the issues with Vserver DR components.

### Syslog Message

Vserver DR check between source Vserver %s and destination Vserver %s found issues with the attribute %s.

### Parameters

**source\_vserver** (STRING): Name of the source Vserver of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

**attribute** (STRING): Name of the attribute that is being checked.

## vsdr.periodic.msids.preserve

### Severity

ALERT

### Description

Vserver DR subsystem has detected that volume MSIDs are not preserved. This would require remounting NFS clients after a DR fail-over. Such relationships may not be supported in future ONTAP releases.

Convert the relationship into an MSID preserving relationship. As part of conversion, the DR Vserver will be stopped and the volume MSIDs will be re-stamped.

### Corrective Action

See the KB article <https://kb.netapp.com/support/s/article/How-to-convert-DR-Relationship-into-MSID-preserving-relation> for more information on how to convert the relationship into an MSID preserving one.

### Syslog Message

Vserver DR: Possibly deprecated variant of Vserver DR relationship for the destination Vserver %s.

### Parameters

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

## vsdr.qos events

### vsdr.QoS.MaxQoSPGs

#### Severity

ERROR

#### Description

This message occurs during QoS configuration replication in a MetroCluster(tm) setup when there are maximum number of QoS policy groups already existing on the destination cluster and the system attempts to create a new QoS policy group.

#### Corrective Action

Delete few QoS policy groups on the destination cluster using "qos policy-group delete" command and then resynchronize the configuration from the source cluster using "metrocluster vserver resync" command.

#### Syslog Message

Reached the maximum of %d QoS policy groups. Cannot create QoS policy group %s in Vserver %s.

#### Parameters

**max\_QoS\_groups** (INT): Maximum number of QoS policy groups.

**QoS\_policy\_group** (STRING): QoS policy group name.

**vserver\_name** (STRING): Vserver name.

### vsdr.qos.pg.conflict

#### Severity

NOTICE

#### Description

This message occurs during the configuration replication phase of a Vserver DR transfer when the Quality of Service (QoS) policy group that the system is trying to replicate already exists in the cluster on a different Vserver.

#### Corrective Action

Remove the existing QoS policy group or make the Vserver that already has the QoS policy group in it the destination in the Vserver DR relationship.

## Syslog Message

QoS policy group %s cannot be created in Vserver %s as the same QoS policy exists in Vserver %s.

## Parameters

**qos\_policy\_group** (STRING): Name of the QoS policy group that the system is trying to replicate.

**destination\_vserver** (STRING): Name of the destination Vserver into which the QoS policy group is being replicated.

**found\_in\_vserver** (STRING): Name of the Vserver in which the same QoS policy group was found.

# vsdr.rootvol events

## vsdr.rootvol.has.data

### Severity

NOTICE

### Description

This message occurs when the system detects that there is data in the Server Virtual Machine root volume during a disaster recovery (DR) transfer phase. Vserver DR does not protect root volumes.

### Corrective Action

Move the data to a nonroot volume of the Vserver if it needs to be protected using DR.

## Syslog Message

Vserver %s has data in the root volume. Vserver DR does not protect data in the root volume.

## Parameters

**vserver** (STRING): Name of the Vserver.

# vsdr.ruleaddedtosmpolicy events

## vsdr.RuleAddedToSmPolicy

### Severity

NOTICE

### Description

This message occurs when a MAX Data® schedule update adds a new rule to a SnapMirror® policy with the "mirror-vault" type. The update is part of an SVM disaster recovery (SVMDR) relationship.

### Corrective Action

(None).

## Syslog Message

New rule with SnapMirror label "%s" is added to SnapMirror policy "%s" of Vserver "%s" used by the Vserver DR relationship.

## Parameters

**snapmirror\_label** (STRING): Name of the SnapMirror label.

**sm\_policy\_name** (STRING): Name of the SnapMirror policy.

**vserver\_name** (STRING): Name of the Vserver of Vserver DR relationship.

## vsdr.rulemodifiedofsmpolicy events

### vsdr.RuleModifiedOfSmPolicy

#### Severity

ALERT

#### Description

This message occurs when a rule attached to a SnapMirror® policy with "mirror-vault" type must be modified, so that L2 scheduled snapshots can be reused during a MAX Data schedule update.

#### Corrective Action

Modify the SnapMirror policy rule by using the "snapmirror policy modify-rule -vserver <vserver-name> -policy <policy-name> -snapmirror-label<snapmirror-label-name> -preserve false" command.

#### Syslog Message

Rule with SnapMirror label "%s" present in SnapMirror policy "%s" on Vserver "%s" used by the Vserver DR relationship needs to be modified.

#### Parameters

**snapmirror\_label** (STRING): Name of the SnapMirror label.

**sm\_policy\_name** (STRING): Name of the SnapMirror policy.

**vserver\_name** (STRING): Name of the Vserver of Vserver DR relationship.

## vsdr.schedule events

### vsdr.schedule.unavailable

#### Severity

ERROR

#### Description

This message occurs when a schedule is not available on the secondary site for replication of policies for a Vserver in a MetroCluster(tm) configuration or Vserver DR setup.

#### Corrective Action

Create a new job schedule on the remote site that is identical to that on the source site by using the 'job schedule cron create' command. If this is a SnapMirror® or MetroCluster(tm) environment, run the "snapmirror resync" command or the "metrocluster vserver resync" command respectively to resynchronize the Vserver configuration.

#### Syslog Message

Configuration replication for policy "%s" in Vserver "%s" has failed because the schedule mentioned in the policy is not available.

#### Parameters

**policy\_name** (STRING): Name of the policy that has a schedule that is unavailable on the secondary site.

**vserver\_name** (STRING): Vserver name on which the policy was associated.

# vsdr.vol events

## vsdr.vol.aff.inline.dedup

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem does not enable inline deduplication on the destination volume. Inline deduplication is not enabled when the destination volume is on neither an AFF system nor a hybrid aggregate.

### Corrective Action

If inline deduplication is required on the destination volume, move the volume to a hybrid aggregate, and then enable inline deduplication. For more information, refer to TR-4476.

### Syslog Message

Vserver DR: The volume %s of the Vserver %s is on neither an AFF system nor a hybrid aggregate. Therefore, inline deduplication is not enabled on the volume.

### Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.  
**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

## vsdr.vol.aggregateoffline

### Severity

ALERT

### Description

This message occurs during configuration replication, when the aggregate on which the volume needs to be created or modified is either offline or not in the cluster.

### Corrective Action

Log in to the cluster and bring the aggregate online by using the "storage aggregate online" command. If the node that owns the aggregate is down, bring it back up.

### Syslog Message

Configuration replication for volume "%s" in Vserver "%s" has failed because the aggregate "%s" is offline.

### Parameters

**volume\_name** (STRING): Volume on which this error occurred.  
**aggregate\_name** (STRING): Name of the aggregate.  
**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.vol.clusterlifedown

### Severity

ERROR

## Description

This message occurs during configuration replication, when the the volume creation fails because the aggregate is not accessible because the cluster logical interface (LIF) is down.

## Corrective Action

Log in to the source cluster and run the "network interface show" command to see the state of the cluster LIFs and to identify the down LIFs. Verify that port connectors, cabling, and cluster switch ports are correctly installed and configured.

## Syslog Message

Configuration replication for volume "%s" in Vserver "%s" has failed because the cluster LIF is operationally down.

## Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.vol.compr.mismatch

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem discovers a mismatch between the compression types of the source and destination volumes, and therefore do not apply the compression type setting of the source volume to the destination volume.

### Corrective Action

Undo all off the compression savings on the destination volume, and then modify the compression type of the volume. For more information, refer to TR-4476.

### Syslog Message

Vserver DR: Compression type of destination volume %s of Vserver %s is %s, which is different from the source volume compression type %s. Therefore, the compression type of the source volume was not applied on the destination volume.

### Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

**destination\_compression\_type** (STRING): Compression type of the destination volume.

**source\_compression\_type** (STRING): Compression type of the source volume.

## vsdr.vol.duplicate.mdv.name

### Severity

ERROR

### Description

This message occurs when the configuration of a metadata volume cannot be replicated because there exists another metadata volume with the same name on the destination cluster.

## Corrective Action

Contact NetApp technical support to remove one of the duplicate metadata volumes.

## Syslog Message

Configuration replication for metadata volume "%s" in Vserver "%s" has failed because there exists another metadata volume with the same name on the destination cluster.

## Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Name of the Vserver for which configuration replication is happening.

## vsdr.vol.encryptsrc.plaindst

### Severity

ALERT

### Description

This message occurs when an encrypted volume at the source of a Vserver-DR setup is replicated as a non-encrypted volume at the destination. This can happen if: 1) the key management capability is not enabled at the destination Vserver; 2) the destination Vserver does not have an aggregate that supports encrypted volumes; or 3) the volume of the source Vserver has undergone a conversion operation from non-encrypted to encrypted.

### Corrective Action

If necessary, set up the key management capability at the destination by using the "security key-manager setup" command. Ensure that the destination Vserver contains an aggregate that supports encrypted volumes. Use the "snapmirror show" command to determine the type of relationship. If the displayed type is "XDP", encrypt the destination volume by performing a conversion operation from non-encrypted to encrypted using the "volume encryption conversion start" command. However, if the type is "DP", encrypt the destination volume by performing a volume move operation using the "volume move start -encrypt -destination true" command.

### Syslog Message

Source volume for volume "%s" in Vserver "%s" is encrypted, but it is replicated as a non-encrypted volume at the destination.

### Parameters

**dest\_volume\_name** (STRING): Name of the destination volume.

**source\_vserver\_name** (STRING): Name of the source Vserver.

**dest\_vserver\_name** (STRING): Name of the destination Vserver.

## vsdr.vol.junction.clash

### Severity

ERROR

### Description

This message occurs when the system detects a conflict in the junction path between the volumes of a source and destination Vserver during a SnapMirror resync operation.



## Corrective Action

Unmount the destination volume or mount it at a different path, and then try the operation again.

## Syslog Message

Source Vserver "%s" has volume "%s" and junction path "%s" that conflicts with destination Vserver "%s", volume "%s" and the same junction path "%s". Reason: The destination volume is excluded from Vserver DR protection in the reverse direction.

## Parameters

**source\_vserver\_name** (STRING): Name of the source Vserver.  
**src\_volume\_name** (STRING): Name of the source volume.  
**src\_volume\_junction\_path** (STRING): Junction path of the source volume.  
**dest\_vserver\_name** (STRING): Name of the destination Vserver.  
**dest\_volume\_name** (STRING): Name of the destination volume.  
**dest\_volume\_junction\_path** (STRING): Junction path of the destination volume.

## vsdr.vol.move.clone.child

### Severity

ERROR

### Description

This message occurs when a clone volume is moved in a Vserver DR setup. As a direct consequence, an implicit clone split occurs. The parent volume no longer has any relation with the child volume. Snapshots(tm) copies of the parent volume are unlocked and can get deleted, thereby affecting the SnapMirror® replication.

### Corrective Action

Run "volume move" command to move the volume to the corresponding aggregate.

### Syslog Message

The volume "%s" of Vserver "%s" has been moved. It is a clone volume. Move the clone volume as well.

### Parameters

**volume\_name** (STRING): Volume that has to be manually moved.  
**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.vol.move.clone.parent

### Severity

ERROR

### Description

This message occurs when the parent of a clone volume is moved in a Vserver disaster-recovery (DR) configuration. If the moved volume is deleted or its Snapshot® copies are deleted, the SnapMirror® transfer fails.

### Corrective Action

Run the "volume move" command to move the volume to an appropriate aggregate. If the volume was deleted at the source and the transfer failed, additionally run "snapmirror resync".

## Syslog Message

The volume "%s" of Vserver "%s" has been moved. The mirroring of this volume is at risk as it is the parent of a clone volume.

## Parameters

**volume\_name** (STRING): Volume that has to be manually moved.

**vserver\_name** (STRING): Name of the Vserver for which configuration replication is happening.

## vsdr.vol.name.clash

### Severity

ERROR

### Description

This message occurs when the source Vserver contains a volume with the same name as one on the destination Vserver, and the volume at the destination is excluded from Vserver Disaster Recovery(DR) Protection in the reverse direction. Future resynchronization attempts might fail unless the problem is resolved.

### Corrective Action

Rename the source or the destination volume, and then try the operation again.

## Syslog Message

The source Vserver "%s" has a volume with name "%s" while the destination Vserver "%s" has a volume with the same name "%s" that is excluded from Vserver DR Protection in the reverse direction.

## Parameters

**source\_vserver\_name** (STRING): Name of the source Vserver in a Vserver DR setup.

**src\_volume\_name** (STRING): Name of the volume on the source Vserver.

**dest\_vserver\_name** (STRING): Name of the destination Vserver in a Vserver DR setup.

**dest\_volume\_name** (STRING): Name of the volume on the destination Vserver.

## vsdr.vol.noparentvolume

### Severity

ERROR

### Description

This message occurs when the volume that needs to be created on the destination cluster has a parent volume that did not get replicated. The possible types of volumes for which this error can occur are LS, FlexCache®, and FlexClone®.

### Corrective Action

Log in to the cluster and unlock the Vserver by using "vserver unlock" command at a diag privilege level, and then create the parent volume by using the "volume create" command. Lock the Vserver by using the "vserver lock" command, and start the configuration replication with a resync.

## Syslog Message

Configuration replication for volume "%s" of type %s in Vserver "%s" has failed because the parent volume for this volume has not been replicated.

## Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**volume\_type** (STRING): Type of volume. Possible values could be LS, FlexCache, or FlexClone.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.vol.nospaceinaggregate

### Severity

ERROR

### Description

This message occurs when there is no space left in the aggregate for new volumes to be created during configuration replication.

### Corrective Action

Log in to the cluster and assign more disks to the aggregate.

### Syslog Message

Configuration replication for volume "%s" in Vserver "%s" has failed because the aggregate "%s" does not have enough space.

## Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

**aggregate\_name** (STRING): Name of the aggregate.

## vsdr.vol.nospaceinfabricpoolaggregate

### Severity

ERROR

### Description

This message occurs when there is not enough space left in the fabric pool aggregate for new volumes to be created during configuration replication.

### Corrective Action

Log in to the cluster and assign more disks to the perf tier of fabric pool aggregate.

### Syslog Message

Couldn't place volume "%s" of Vserver "%s" in fabric pool aggregate as there is not enough spece left in perf tier.

## Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.vol.protect.childmnt

### Severity

ERROR

## Description

This message occurs when a protected volume is mounted under an unprotected volume on the source of a Storage Virtual Machine (SVM) disaster recovery (DR) relationship and a SnapMirror® transfer is subsequently initiated. Data continues to be served from this volume, but the volume is not protected by the DR relationship.

## Corrective Action

Execute any of the following operations on the source cluster to ensure that parent and child volumes have matching protection types, and then invoke the "snapmirror resync" command. 1. Mount the volume under a protected volume. 2. Change the protection type of the parent volume in the junction-path to protected. 3. Change the protection type of the volume to unprotected. To change the protection type, use the command "volume modify -vserver vserver\_name -volume volume\_name -vserver-dr-protection".

## Syslog Message

Unable to mount the volume %s:%s under an unprotected volume.

## Parameters

**vserver\_name** (STRING): Name of the SVM for which the SnapMirror transfer was initiated.

**volume\_name** (STRING): Name of the volume that the system is trying to replicate on the destination SVM.

## vsdr.vol.quota.off.fail

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem fails to disable quotas on the destination volume.

### Corrective Action

Disable quotas on the volume by using the "volume quota off".

### Syslog Message

Vserver DR: Failed to disable quotas on volume %s of Vserver %s. Reason: %s

### Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

**reason** (STRING): Reason.

## vsdr.vol.quota.on.fail

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem fails to enable quotas on the destination volume.

### Corrective Action

Enable quotas on the volume by using the "volume quota on" command.

## Syslog Message

Vserver DR: Failed to enable quotas on volume %s of Vserver %s. Reason: %s

## Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

**reason** (STRING): Reason.

## vsdr.vol.sg.data.compact

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem does not enable data compaction on the destination volume. Data compaction is not enabled when the destination volume space guarantee setting is not set to "none".

### Corrective Action

Set the space guarantee of the volume to "none", and then enable data compaction. For more information, refer to TR-4476.

## Syslog Message

Vserver DR: Volume space guarantee setting for volume %s of Vserver %s is not set to "none". Therefore, data compaction is not enabled on the volume.

## Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

## vsdr.vol.sg.inline.dedup

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem does not enable cross volume inline deduplication on the destination volume. Inline deduplication is not enabled when the destination volume space guarantee setting is not set to "none".

### Corrective Action

Set the space guarantee of the volume to "none", and then enable cross volume inline deduplication. For more information, refer to TR-4476.

## Syslog Message

Vserver DR: Volume space guarantee setting for volume %s of Vserver %s is not set to "none". Therefore, cross volume inline deduplication is not enabled on the volume.

## Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.

## **vsdr.vol.sis.capture.fail**

### **Severity**

ALERT

### **Description**

This message occurs when a Vserver DR subsystem fails to replicate the efficiency configuration from the source volume to the destination volume.

### **Corrective Action**

Manually enable the efficiency configuration on the destination volume. For more information, refer to TR-4476.

### **Syslog Message**

Vserver DR: Failed to replicate the efficiency configuration on volume %s of Vserver %s. Reason: %s

### **Parameters**

**source\_volume** (STRING): Name of the source volume of the Vserver DR relationship.  
**source\_vserver** (STRING): Name of the source Vserver of the Vserver DR relationship.  
**reason** (STRING): Reason for the failure.

## **vsdr.vol.sis.fail**

### **Severity**

ALERT

### **Description**

This message occurs when a Vserver DR subsystem fails to apply the efficiency configuration on the destination volume.

### **Corrective Action**

Manually apply the efficiency configuration on the volume. For more information, refer to TR-4476.

### **Syslog Message**

Vserver DR: Failed to apply the efficiency configuration on volume %s of Vserver %s. Reason: %s

### **Parameters**

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.  
**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.  
**reason** (STRING): Reason.

## **vsdr.vol.vldboffline**

### **Severity**

ALERT

### **Description**

This message occurs when the volume location database (VLDB) goes offline during configuration replication. None of the volume services can commence without bringing it online.

### Corrective Action

Retry the SnapMirror® operation by running the "snapmirror resync" command. If the problem persists, contact NetApp technical support.

### Syslog Message

Configuration replication for volume "%s" in Vserver "%s" has failed because the VLDB is offline.

### Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.volmov events

### vsdr.volmov.msids.conflict

#### Severity

ERROR

#### Description

This message occurs during a SnapMirror® resync operation on a Storage Virtual Machine (SVM) disaster recovery (DR) relationship, when the MSID of a volume on the source SVM is the same as the MSID of a temporary volume on the destination SVM and the temporary volume is created for a volume move operation. The resync operation fails to start.

#### Corrective Action

Abort the volume move operation by using the "volume move abort" command, or wait for the volume move operation to finish before issuing another "snapmirror resync" command on the SVM DR relationship.

#### Syslog Message

The MSID of temporary volume "%s" on the destination SVM "%s" conflicts with the MSID of volume "%s" on the source.

#### Parameters

**tmp\_volume\_name** (STRING): Volume name of the temporary volume on the destination SVM.

**vserver\_name** (STRING): SVM name for which the SnapMirror resync was initiated.

**src\_volume\_name** (STRING): Volume name on the source SVM.

## vsdr.volmove events

### vsdr.volmove.esrc.plaindst

#### Severity

ALERT

#### Description

This message occurs when a volume at the source of a Vserver-DR setup is encrypted, but the volume at the destination is not configured to be encrypted.

#### Corrective Action

Verify that the intended destination aggregate supports encryption and Key Management capability is

enabled. If so, make the destination volume encrypted by performing a volume move operation at the destination: "volume move start -vserver <vserver-name> -volume <volume-name> -destination-aggregate <aggregate> -encrypt-destination true".

### Syslog Message

Source volume for volume "%s" in Vserver "%s" is encrypted, but the destination is not.

### Parameters

**dest\_vserver\_name** (STRING): Name of the destination Vserver.

**dest\_volume\_name** (STRING): Name of the destination volume.

## vsdr.volmovsrcfpnonfpaggr events

### vsdr.volMovSrcFPNonFPAggr

#### Severity

NOTICE

#### Description

This message occurs when a volume move operation is executed between FabricPool and non-FabricPool aggregates on the source Vserver of an SVM disaster recovery (SVMDR) relationship.

#### Corrective Action

If the aggregate type hosting the volume differs at the primary and secondary clusters, move the destination Vserver of the SVMDR relationship at the secondary cluster by using the "volume move start -vserver <vserver-name> -volume <volume-name> -destination-aggregate <aggregate>" command.

#### Syslog Message

Source volume "%s" on Vserver "%s" of the SVMDR relationship is moved from "%s" aggregate "%s" to "%s" aggregate "%s".

#### Parameters

**source\_volume\_name** (STRING): Name of the source volume of the SVMDR relationship.

**source\_vserver\_name** (STRING): Name of the source Vserver of the SVMDR relationship.

**src\_aggr\_type** (STRING): Source aggregate type, either FabricPool or non-FabricPool.

**source\_aggr\_name** (STRING): Name of the source aggregate at the primary cluster.

**dest\_aggr\_type** (STRING): Destination aggregate type, either FabricPool or non-FabricPool.

**dest\_aggr\_name** (STRING): Name of the destination aggregate at the primary cluster.

## vsdr.volume events

### vsdr.volume.idconflict

#### Severity

ERROR

#### Description

This message occurs when the the volume to be created on the destination cluster fails because of an internal identity conflict with another volume in the system.



### Corrective Action

Log in to the source cluster and delete the volume by using the "volume delete" command.

### Syslog Message

Configuration replication for volume "%s" in Vserver "%s" has failed because of an internal ID conflict.

### Parameters

**volume\_name** (STRING): Volume on which this error occurred.

**vserver\_name** (STRING): Vserver name for which configuration replication is happening.

## vsdr.volume.name.conflict

### Severity

ERROR

### Description

This message occurs during a SnapMirror resync® operation on a Storage Virtual Machine (SVM) disaster recovery (DR) relationship, when a new volume needs to be created on the destination SVM but another volume with the same name already exists on it, and it is also the destination of a volume-level SnapMirror relationship. The resync operation fails to start.

### Corrective Action

Rename the volume on the destination SVM by using the "volume rename" command or delete it by using the "volume delete" command, and then execute another "snapmirror resync" command.

### Syslog Message

The volume "%s" is already present in SVM "%s" and it is the destination of a volume-level SnapMirror relationship.

### Parameters

**volume\_name** (STRING): Name of the volume that the system is trying to replicate into the destination SVM.

**vserver\_name** (STRING): SVM name for which the SnapMirror resync was initiated.

## vsdr.volume.sis.on.fail

### Severity

ALERT

### Description

This message occurs when a Vserver DR subsystem fails to enable volume efficiency on the volume.

### Corrective Action

Enable volume efficiency on the destination volume by using the "volume efficiency on" command. For more information, refer to TR-4476.

### Syslog Message

Vserver DR: Volume efficiency configuration is not applied on volume %s of Vserver %s. Reason: %s

### Parameters

**destination\_volume** (STRING): Name of the destination volume of the Vserver DR relationship.

**destination\_vserver** (STRING): Name of the destination Vserver of the Vserver DR relationship.  
**reason** (STRING): Reason.

# vserver events

## vserver.config events

### vserver.config.createDirectoryFailed

**Severity**

ERROR

**Description**

This message occurs when the Vserver create operation cannot create the Vserver-specific configuration directory.

**Corrective Action**

Ensure that, the root volume File System(FS) is available and has enough space to create Vserver-specific configuration directory.

**Syslog Message**

"Failed to create the Vserver-specific configuration directory for Vserver %s. Reason: The root volume is full or not available."

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

### vserver.config.nameService.createDirectoryFailed

**Severity**

ERROR

**Description**

This message occurs when Vserver-specific name-service configuration directory cannot be created.

**Corrective Action**

Ensure that, the root volume File System(FS) is available and has enough space to create Vserver-specific name-service configuration directory.

**Syslog Message**

Failed to create the Vserver-specific name-service configuration directory for Vserver %s. Reason: The root volume is full or not available.

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

## vserver.ipspace events

### vserver.ipspace.lookupfailed

**Severity**

ALERT

## Description

This message occurs in a MetroCluster® configuration when attempting to replicate a "vserver create" operation onto the destination cluster, but the configured IPspace cannot be found on the destination cluster. This error can occur due to networking subsystem issues on the destination cluster.

## Corrective Action

The operation is retried several times. Use the "metrocluster vserver show" command to verify that the configuration state of the Vserver becomes healthy. If the configuration state does not become healthy after an hour, contact NetApp technical support for assistance.

## Syslog Message

Vserver "%s" cannot be created because IPspace information related to the Vserver is not currently available.

## Parameters

**vserver\_name** (STRING): Name of the Vserver that is being replicated.

## vserver.ipspace.notAvailable

### Severity

ERROR

### Description

This message occurs in a MetroCluster® configuration when attempting to replicate a "vserver create" operation onto the destination cluster, but the IPspace name for creation of the Vserver has not yet propagated to the destination cluster. This error can occur due to networking subsystem issues on the source cluster.

### Corrective Action

The operation is retried several times. Use the "metrocluster vserver show" command to verify that the configuration state of the Vserver becomes healthy. If the configuration state does not become healthy after an hour, contact NetApp technical support for assistance.

### Syslog Message

Vserver "%s" cannot be created because IPspace information related to the Vserver is not currently available.

### Parameters

**vserver\_name** (STRING): Name of the Vserver that is being replicated.

## vserver.ipspace.notfound

### Severity

ERROR

### Description

This message occurs in a Metrocluster configuration when we are attempting to replicate a Vserver create operation on to the destination cluster and the Ipspace provided does not exist on the destination cluster.

### Corrective Action

Create an Ipspace with the same name as that of the ipspace used for creating the Vserver on the source cluster.

## Syslog Message

Vserver %s cannot be created as ipspace %s does not exist.

## Parameters

**vserver\_name** (STRING): Name of the vsriver we are trying to replicate.

**ipspace\_name** (STRING): Name of the ipspace on which we are trying to create the Vserver.

# vserver.mcc events

## vserver.mcc.repl.failed

### Severity

ERROR

### Description

This message occurs in a MetroCluster™ configuration when an attempt to create a Vserver on the partner cluster fails.

### Corrective Action

Run the 'metrocluster vsriver show' command, and follow the corrective action specified in the output.

## Syslog Message

Vserver %s cannot be created on the partner cluster.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

## vserver.mcc.start.failed

### Deprecated

Deprecated as of version 9.2.0 as this EMS is not required.

### Severity

ALERT

### Description

This message occurs when a sync-destination Vserver fails to get started during a MetroCluster(tm) switchover operation or a sync-source Vserver fails to get started during a MetroCluster switchover operation.

### Corrective Action

Start the Vserver which failed to get started during MetroCluster switchover/switchover.

## Syslog Message

Failed to start the Vserver %s during Metrocluster switchover/switchover operation.

## Parameters

**vserver\_name** (STRING): Name of the vsriver which failed to get started.

## **vserver.mcc.stop.failed**

### **Deprecated**

Deprecated as of version 9.2.0 as this EMS is not required.

### **Severity**

ERROR

### **Description**

This message occurs when a sync-destination Vserver fails to get stopped during a MetroCluster switchback operation.

### **Corrective Action**

Stop the Vserver which failed to get stopped during MetroCluster switchback.

### **Syslog Message**

Failed to stop the Vserver %s during MetroCluster switchback operation.

### **Parameters**

**vserver\_name** (STRING): Name of the vserver which failed to get stopped during MetroCluster switchback.

## **vserver.name events**

### **vserver.name.conflict.local**

#### **Severity**

NOTICE

#### **Description**

This message occurs when use of the "vserver rename" command creates a Vserver name conflict in the local cluster, and the local cluster resolves the name conflict by incorporating the Vserver Unique Identifier (UID) into the name. This creates a new and unique Vserver name, but that new name is not the name specified in the original rename command.

#### **Corrective Action**

None

#### **Syslog Message**

Vserver name conflict in local cluster with existing local or peer Vserver %s. Local cluster has assigned unique Vserver name %s

#### **Parameters**

**vserver\_name** (STRING): Vserver name that caused the renaming conflict.

**unique\_vserver\_name** (STRING): Unique Vserver name that Local cluster has assigned.

### **vserver.name.conflict.peer**

#### **Severity**

NOTICE

## Description

This message occurs when use of the "vserver rename" command creates a Vserver name conflict in the peer cluster, and the peer cluster resolves the name conflict by incorporating the Vserver Unique Identifier (UID) into the name. This creates a new and unique Vserver name, but that new name is not the name specified in the original rename command.

## Corrective Action

None

## Syslog Message

Vserver name conflict occurred in peer cluster %s with existing local or peer Vserver %s. Peer cluster has assigned unique Vserver name %s.

## Parameters

**peer\_cluster** (STRING): Peer cluster name.

**vserver\_name** (STRING): Vserver name that caused the renaming conflict.

**unique\_vsname** (STRING): Unique Vserver name that Peer cluster has assigned.

# vserver.oper events

## vserver.oper.state.down

### Severity

ERROR

### Description

This message occurs when the system detects a cluster reboot on a Disaster Recovery Source (MetroCluster(tm) configuration or Vserver DR). All Vservers that are a source of the disaster-recovery relationship are brought down operationally if the destination is activated. This is done to prevent an identity clash in case the Secondary site is brought up.

### Corrective Action

If the destination Vservers or site are already brought up, nothing needs to be done. Otherwise, if the source is brought down, then unlock the Vservers by using the "vserver unlock" command, and then bring them up operationally by using the "vserver start" command.

### Syslog Message

The Vservers on cluster %s that are in a DR relationship have been operationally brought down.

### Parameters

**cluster\_name** (STRING): Name of the source cluster.

# vserver.peering events

## vserver.peering.relCreated

### Severity

NOTICE

**Description**

This message occurs when the Vserver peer relationship is established.

**Corrective Action**

(None).

**Syslog Message**

Vserver peer relationship between %s and %s is established having remote peer Vserver name %s.

**Parameters**

**vserver** (STRING): Name of the Vserver for which a peer is created.

**peer\_vserver\_localname** (STRING): Name locally used for the peer Vserver created.

**peer\_vserver\_remotename** (STRING): Actual name of the peer Vserver.

**vserver.peering.relDeleted****Severity**

INFORMATIONAL

**Description**

This message occurs when the Vserver peer relationship is deleted successfully.

**Corrective Action**

(None).

**Syslog Message**

"Vserver peer relationship between %s and %s is deleted."

**Parameters**

**vserver** (STRING): Name of the Vserver.

**peer\_vserver** (STRING): Name of the peer Vserver.

**vserver.peering.relResumed****Severity**

INFORMATIONAL

**Description**

This message occurs when the Vserver peer relationship is resumed successfully.

**Corrective Action**

(None).

**Syslog Message**

"Vserver peer relationship between %s and %s is resumed."

**Parameters**

**vserver** (STRING): Name of the Vserver.

**peer\_vserver** (STRING): Name of the peer Vserver.



## **vserver.peering.relSuspended**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the Vserver peer relationship is suspended successfully.

### **Corrective Action**

(None).

### **Syslog Message**

"Vserver peer relationship between %s and %s is suspended."

### **Parameters**

**vserver** (STRING): Name of the Vserver.

**peer\_vserver** (STRING): Name of the peer Vserver.

## **vserver.peering.remCallFailed**

### **Severity**

ERROR

### **Description**

This message occurs when any Vserver peering command fails because of a communication failure with a remote system.

### **Corrective Action**

Use the "network interface show" command to verify that intercluster logical interfaces (LIFs) on both clusters are up. Use the "cluster peer health show -bypass-cache true" command to verify that cluster peering is healthy. If you want to retry a failed delete, suspend, or resume operation, consider using the "force" option, which bypasses network connectivity issues.

### **Syslog Message**

"Vserver peering command '%s' failed. Reason: %s. Vserver peering states might be inconsistent."

### **Parameters**

**command** (STRING): Vserver peering command.

**reason** (STRING): Reason for failure.

## **vserver.rbac events**

### **vserver.rbac.alias.command**

#### **Deprecated**

ONTAP now directly creates a role entry on the target of the alias.

### **Severity**

ERROR

## Description

This message occurs when a custom role associated with a Storage Virtual Machine (SVM) has an aliased command directory associated with it. SVM custom roles must be associated with actual command directories and commands, not aliases to them. The aliased commands associated with SVM custom roles will not be available under SVM context.

## Corrective Action

Delete the role on the SVM, and then re-create it using the actual command.

## Syslog Message

The custom role %s is associated with alias command - %s on SVM %s. Create role with actual command - %s.

## Parameters

**role\_name** (STRING): Name of the role for which creation failed.  
**alias\_name** (STRING): Name of the alias that is associated with the role.  
**vserver\_name** (STRING): Name of the SVM.  
**non\_alias\_name** (STRING): Name of the actual command.

## vserver.rbac.replicateRole

### Severity

ERROR

### Description

This message occurs when replication of a role is failed on a Vserver because the command directory associated with it is an alias.

### Corrective Action

Delete the role on the source Vserver, and then re-create it using the non-alias target. Then retry the operation again.

### Syslog Message

Failed to replicate the role %s associated with the alias %s on Vserver %s.

### Parameters

**role\_name** (STRING): Name of the role for which creation failed.  
**alias\_name** (STRING): Name of the alias that is associated with the role.  
**vserver\_name** (STRING): Name of the Vserver.

## vserver.rename events

### vserver.rename.fail.peer

#### Severity

ERROR

#### Description

This message occurs when the "vserver rename" command fails to update the peer Vserver name on the peer cluster in a Vserver peer relationship.

### **Corrective Action**

Use the (privilege: advanced) "vserver peer repair-peer-name" command in the local cluster to update the peer Vserver name in the Vserver peer relationships present on the peer cluster.

### **Syslog Message**

Failed to update new peer Vserver %s name on peer cluster %s.

### **Parameters**

**vserver\_name** (STRING): Vserver name that caused the update attempt to fail.

**peer\_cluster** (STRING): Peer cluster name.

## **vserver.rootvolume events**

### **vserver.rootVolume.full**

#### **Severity**

ERROR

#### **Description**

This message occurs when there is no space left on Vserver root Volume.

#### **Corrective Action**

Increase the size of the root volume. To increase the size of the volume, run the 'volume modify' command.

#### **Syslog Message**

Vserver %s has its root volume full.

#### **Parameters**

**vserver\_name** (STRING): Name of the vserver with root volume full.

### **vserver.rootVolume.noInodes**

#### **Severity**

ERROR

#### **Description**

This message occurs when there are no inodes are available on Vserver root volume.

#### **Corrective Action**

Increase the size of the root volume. To increase the size of the volume, run the 'volume modify' command.

#### **Syslog Message**

Vserver %s has no inodes in its root volume.

#### **Parameters**

**vserver\_name** (STRING): Name of the vserver with no inodes left in root volume.

### **vserver.rootVolume.offline**

**Severity**

ERROR

**Description**

This message occurs when the Vserver root Volume is offline or not accessible for read/write operations.

**Corrective Action**

Make the vservers root volume online.

**Syslog Message**

Vserver %s has its root volume offline.

**Parameters**

**vserver\_name** (STRING): Name of the vservers for which the root volume is not accessible.

## vserver.start events

### vserver.start.failed

**Severity**

ERROR

**Description**

This message occurs when a 'vserver start' operation fails due to failure to start either a logical interface (LIF) or a protocol for the Vserver specified.

**Corrective Action**

Ensure that the network configuration, LIF details, and protocol configurations are correct, then try to start the Vserver again.

**Syslog Message**

Failed in Vserver start due to failure to start either a LIF or a protocol for the Vserver %s (UUID: %s).

**Parameters**

**vserver\_name** (STRING): Name of the Vserver.

**vserver\_uuid** (STRING): UUID of the Vserver.

### vserver.start.succeeded

**Severity**

INFORMATIONAL

**Description**

This message occurs when a 'vserver start' operation succeeds.

**Corrective Action**

(None).

**Syslog Message**

Vserver %s (UUID: %s) started successfully.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**vserver\_uuid** (STRING): UUID of the Vserver.

# vserver.stop events

## vserver.stop.failed

### Severity

ERROR

### Description

This message occurs when a 'vserver stop' operation fails due to failure to stop either a logical interface (LIF) or a protocol for the Vserver specified.

### Corrective Action

Ensure that the network configuration, LIF details, and protocol configurations are correct, then try to stop the Vserver again.

### Syslog Message

Failed in Vserver stop due to failure to stop either a LIF or a protocol for the Vserver %s (UUID: %s).

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**vserver\_uuid** (STRING): UUID of the Vserver.

## vserver.stop.succeeded

### Severity

INFORMATIONAL

### Description

This message occurs when a 'vserver stop' operation succeeds.

### Corrective Action

(None).

### Syslog Message

Vserver %s (UUID: %s) stopped successfully.

## Parameters

**vserver\_name** (STRING): Name of the Vserver.

**vserver\_uuid** (STRING): UUID of the Vserver.

# vserver.storage events

## vserver.storage.full

**Severity**

ERROR

**Description**

This message occurs when the Storage of Vserver is full.

**Corrective Action**

Increase the available size of Vserver either by raising storage-limit or by downsizing or deleting volumes.

**Syslog Message**

Vserver "%s" is full by consuming the storage limit "%s".

**Parameters**

**vserver\_name** (STRING): Vserver name for which the storage is full.

**storage\_limit** (STRING): The storage-limit for the Vserver.

**vserver.storage.nearlyFull****Severity**

ALERT

**Description**

This message occurs when the Vserver storage size exceeds 95% of the storage limit.

**Corrective Action**

Increase the available size of Vserver either by raising storage-limit or by downsizing or deleting volumes.

**Syslog Message**

Vserver "%s" with current storage size of "%s" has exceeded "95%%" of storage limit "%s".

**Parameters**

**vserver\_name** (STRING): Vserver name for which the storage size has exceeded 95% of the storage limit.

**storage\_size** (STRING): The storage used size of the Vserver.

**storage\_limit** (STRING): The storage-limit for the Vserver.

**vserver.storage.size.breached****Severity**

NOTICE

**Description**

This message occurs when tenant has exceeded the capacity quota allotted to the Vserver. The checker automatically deselects this mismatch and attempts a fix. To see details including the updated size use "vserver show -fields storage-limit,storage-allocated,storage-reserved" command.

**Corrective Action**

(None).

**Syslog Message**

Vserver "%s" size "%s" has exceeded storage capacity quota of "%s".

## Parameters

**vserver\_name** (STRING): The Vserver whose size exceeded the capacity quota allotted to the tenant.

**vserver\_size** (STRING): Vserver's size allotted up to a point to the tenant.

**capacity\_limit** (STRING): Capacity limit allotted to the Vserver.

## vserver.storage.size.mismatch

### Severity

NOTICE

### Description

This message occurs when the Vserver's size allotted to the tenant up to a point does not match sum of all the volume's provisioned size that can be used to enforce Vserver capacity quota. The checker automatically deselects this mismatch and attempts a fix. To see details including the updated size use "vserver show -fields storage-limit,storage-allocated,storage-reserved" command.

### Corrective Action

(None).

### Syslog Message

Vserver "%s" size allocated to the tenant up to a point "%s" does not match sum of all the volume's provisioned size "%s".

## Parameters

**vserver\_name** (STRING): Vserver for which the size allotted to the tenant up to a point does not match sum of all the volume's provisioned size.

**vserver\_size** (STRING): Vserver's size allotted up to a point to the tenant.

**provisioned\_size** (STRING): Sum of all the volume's provisioned size in Vserver.

## vserver.storage.threshold

### Severity

ALERT

### Description

This message occurs when the Vserver storage size exceeds the storage-limit-threshold-alert value.

### Corrective Action

Increase the available size of Vserver either by raising storage-limit or by downsizing or deleting volumes.

### Syslog Message

Vserver "%s" with current storage size of "%s" has exceeded "%s%%" of storage limit "%s".

## Parameters

**vserver\_name** (STRING): Vserver name for which the storage size has exceeded the storage-limit-threshold-alert value.

**storage\_size** (STRING): The storage used size of the Vserver.

**storage\_limit\_threshold** (STRING): The storage-limit-threshold-alert value of the Vserver.

**storage\_limit** (STRING): The storage-limit for the Vserver.

# vserver.vsdrr events

## vserver.vsdrr.adt.romode

### Severity

ERROR

### Description

This message occurs in Async Vserver-DR setup, when an auditing-configured Vserver is starting up in Read-Only (RO) mode.

### Corrective Action

None.

### Syslog Message

The Vserver %s is in Async Vserver-DR (RO) mode. Therefore, auditing cannot happen on this Vserver.

### Parameters

**vserver\_name** (STRING): Name of the Vserver that is in Async Vserver-DR (RO) mode.

## vserver.vsdrr.adt.supported

### Severity

INFORMATIONAL

### Description

This message occurs during asynchronous Vserver disaster recovery setup, when the status of Vserver-Disaster Recovery Read-Only mode of a Vserver configured for auditing is set to false when a SnapMirror® relationship breaks.

### Corrective Action

None.

### Syslog Message

Vserver %s is configured with auditing support in Async Vserver-Disaster Recovery setup.

### Parameters

**vserver\_name** (STRING): Name of the Vserver that is in Async Vserver-Disaster Recovery Read only mode.



# wackiron events

## wackiron.near events

### wackiron.near.hour.limit

#### Severity

ERROR

#### Description

This message occurs when the number of hours since wafiron was last run on the specified volume is approaching its limit. If the limit is passed and the bootarg 'wackiron\_enforce\_panic' is true, the system will panic on the next boot.

#### Corrective Action

Run wafiron on the affected volume. From the boot menu, run 'wafiron'. From the dblade CLI, run 'aggr wafiron start [aggr-name]'.

#### Syslog Message

%s: Hours since wafiron was last run is approaching its limit. Hours remaining: %d.

#### Parameters

**volname** (STRING): Name of the aggregate.  
**remaining** (INT): Number of hours remaining.

### wackiron.near.kernel.limit

#### Severity

ERROR

#### Description

This message occurs when the number of different Data ONTAP® kernel versions used since wafiron was last run on the specified volume is approaching its limit. If the limit is passed and the bootarg 'wackiron\_enforce\_panic' is true, the system will panic on the next boot.

#### Corrective Action

Run wafiron on the affected volume. From the boot menu, run 'wafiron'. From the dblade CLI, run 'aggr wafiron start [aggr-name]'.

#### Syslog Message

%s: Number of different Data ONTAP kernel versions used since wafiron was last run is approaching its limit. Data ONTAP kernel version changes remaining: %d.

#### Parameters

**volname** (STRING): Name of the aggregate.  
**remaining** (INT): Number of Data ONTAP kernel version changes remaining.

# wackiron.panic events

## wackiron.panic.warning

### Severity

ERROR

### Description

This message occurs when the bootarg 'wackiron\_enforce\_panic' will cause the system to panic after the next boot if you do not run wafiron on the affected volume before the completion of the next boot.

### Corrective Action

Run wafiron on the affected volume. From the boot menu, run 'wafiron'. From the dblade CLI, run 'aggr wafiron start [aggr-name]'.

### Syslog Message

Warning: System will panic if wafiron is not run on '%s' before the next boot is complete.

### Parameters

**volname** (STRING): Name of the aggregate.

# wackiron.past events

## wackiron.past.hour.limit

### Severity

ERROR

### Description

This message occurs when the number of hours since wafiron was last run on the specified volume is past its limit. Because the limit was passed, if the bootarg 'wackiron\_enforce\_panic' is true, the system will panic on the next boot.

### Corrective Action

Run wafiron on the affected volume. From the boot menu, run 'wafiron'. From the dblade CLI, run 'aggr wafiron start [aggr-name]'.

### Syslog Message

%s: Number of hours since wafiron was last run on the specified volume is past its limit. Run 'wafiron' as soon as possible.

### Parameters

**volname** (STRING): Name of the aggregate.

## wackiron.past.kernel.limit

### Severity

ERROR

**Description**

This message occurs when the number of Data ONTAP® kernel versions used since wafiron was last run on the specified volume is past its limit. Because the limit was passed, if the bootarg 'wackiron\_enforce\_panic' is true, the system will panic on the next boot.

**Corrective Action**

Run wafiron on the affected volume. From the boot menu, run 'wafiron'. From the dblade CLI, run 'aggr wafiron start [aggr-name]'.

**Syslog Message**

%s: Number of Data ONTAP kernel versions used since wafiron was last run on the specified volume is past its limit. Run 'wafiron' as soon as possible.

**Parameters**

**volname** (STRING): Name of the aggregate.

# wafl events

## wafl.32bit events

### wafl.32bit.formatUnsupported

#### Severity

ERROR

#### Description

This message occurs when an aggregate or volume with a 32-bit or mixed format file system is found. These formats are unsupported in this version of Data ONTAP®, and the volume cannot be brought online.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

%s %s%s%s%s has an unsupported 32-bit or mixed format file system (Snapshot copy ID %u).

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

## wafl.8 events

### wafl.8.0.mirror.upgrade.fail

#### Severity

ERROR

#### Description

This message occurs when WAFL® unexpectedly encounters a Data ONTAP® 8.0.x logical mirror volume during an upgrade. This logical mirror volume should have been deleted prior to the upgrade.

#### Corrective Action

Avoid this event by destroying all Data ONTAP 8.0.x logical mirror volumes prior to the upgrade. This event resulted in the volume being marked offline, corrupt, and unrecoverable. Because the current version of Data ONTAP does not support logical mirror volume, the volume cannot be recovered by file system repair mechanisms such as wafiron and cannot be brought online.

#### Syslog Message

Data ONTAP 8.0.x mirror (type %d) cannot be upgraded by the current version of Data ONTAP. Logical mirror volume '%s' was marked offline, corrupt, and unrecoverable. Read the upgrade notes for this release. Such volumes should be deleted prior to the upgrade.

## Parameters

**type** (INT): Mirror type in Data ONTAP 8.0.x.

**vol** (STRING): Volume name.

# wafl.access.events

## wafl.access.fenced.file

### Severity

NOTICE

### Description

This message occurs when the user attempts to access a fenced file. Depending on the type of fence, reads and/or writes will not be permitted for the file.

### Corrective Action

(None).

### Syslog Message

The file with inode %d on volume %s cannot be accessed because it is fenced by %s operation.

## Parameters

**fid** (INT): Inode number of the fenced file.

**vol** (STRING): Name of the volume that contains the fenced file.

**owner** (STRING): Name of the operation that fenced the file.

# wafl.afs.events

## wafl.afs.first.disabled

### Severity

NOTICE

### Description

This message occurs when the deswizzle scanner fails to enable the wafl.deswizzle.afs.first feature on the aggregate because block reallocation redirection cannot be enabled on the aggregate. The deswizzle scanner will run without the wafl.deswizzle.afs.first feature for any volume in the aggregate.

### Corrective Action

(None).

### Syslog Message

afs.first option cannot be enabled on aggregate %s, for volume %s because it is not possible to enable block reallocation on the aggregate.

## Parameters

**aggr** (STRING): Aggregate name.

**vol** (STRING): Volume name.

# waf1.aggr events

## waf1.aggr.addup.offline

### Severity

ERROR

### Description

This message occurs when an aggregate is taken offline due to the aggregate has savings because of aggregate deduplication and this version of ONTAP does not support aggregate level deduplication.

### Corrective Action

Confirm that a version of Data ONTAP which supports aggregate level deduplication is in use.

### Syslog Message

Aggregate name:'%s' has aggregate level deduplication enabled. It is not supported in current Data ONTAP version and was taken offline.

### Parameters

**name** (STRING): Aggregate name.

## waf1.aggr.ar1.callbacks

### Severity

NOTICE

### Description

This message occurs during an aggregate relocation operation, when the system ignores outstanding oplock recalls so that the relocation can progress.

### Corrective Action

(None).

### Syslog Message

Ignoring %d locks during aggregate relocation of %s %s%s%s%s.

### Parameters

**count** (INT): Number of ignored locks.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the aggregate.

**aggr** (STRING): Name of the aggregate.

**app** (STRING): Application UUID.

**aggrident** (STRING): Unique identifier of the aggregate in case the aggregate name itself is insufficient.

## waf1.aggr.btiddb.build

### Severity

INFORMATIONAL

## Description

This message occurs at the completion of an aggregate buftreeid database construction. It records the time spent, inode scanned, msg restart count, and the result.

## Corrective Action

(None).

## Syslog Message

Buftreeid database for aggregate '%s' UUID '%s' was built in %llu msec, after scanning %d inodes and restarting %d times with a final result of %s.

## Parameters

**aggrName** (STRING): Name of the aggregate.

**uuid** (STRING): UUID of the aggregate.

**buildtime** (LONGINT): Total time taken for buftreeid database construction.

**inodes** (INT): Total inodes seen for buftreeid database construction.

**restarts** (INT): Number of message restarts for buftreeid database construction.

**success** (STRING): Whether the constructino was a success or failure.

## wafl.aggr.crt.offline

### Severity

ERROR

### Description

This message occurs when an aggregate is taken offline due to the aggregate was migrated from 7-Mode controller and this version of ONTAP does not support aggregate transition.

### Corrective Action

Confirm that a version of Data ONTAP which supports 7-mode controller migration is in use.

### Syslog Message

Aggregate name:'%s' UUID: '%s' was migrated from 7-Mode controller. It is not supported in current Data ONTAP version and was taken offline.

### Parameters

**name** (STRING): Aggregate name.

**UUID** (STRING): Aggregate UUID.

## wafl.aggr.dump.metadir.rm

### Severity

INFORMATIONAL

### Description

This message occurs when the system cannot remove the dump metadirectory, which was created by a backup job for a mirror. The dump metadirectory is created on the aggregate where the mirror resides, because mirrors are RO volumes. The metadirectory is per-session, which means that the metadirectory is removed at the end of every normal backup. Therefore, the metadirectory will be deleted the next time a backup of the same volume is completed normally.

## Corrective Action

(None).

## Syslog Message

Could not remove %s in aggregate %s for volume %s with UUID %s: %s

## Parameters

**path** (STRING): Path of the dump metadirectory.

**aggr** (STRING): Aggregate where the dump metadirectory resides.

**vol** (STRING): Volume that the dump metadirectory was created for.

**voluuid** (STRING): UUID of the volume that the dump metadirectory was created for.

**err** (STRING): Error indicating why the deletion of the dump metadirectory failed.

## waf1.aggr.giveback.callbacks

### Severity

NOTICE

### Description

This message occurs during a giveback operation, when the system ignores outstanding oplock recalls so that the giveback can progress.

### Corrective Action

(None).

### Syslog Message

Ignoring %d locks when giving back %s %s%s%s%s.

### Parameters

**count** (INT): Number of ignored locks.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the aggregate.

**aggr** (STRING): Name of the aggregate.

**app** (STRING): Application UUID.

**aggrident** (STRING): Unique identifier of the aggregate in case the aggregate name itself is insufficient.

## waf1.aggr.giveback.times

### Severity

INFORMATIONAL

### Description

Issued upon successful completion of giveback of an SFO HA-policy aggregate to record sub-phase timings for different phases of waf1 giveback.

### Corrective Action

(None).

### Syslog Message

(None).



## Parameters

**aggrName** (STRING): The name of the Aggregate.  
**timesString** (STRING): Report of sub phase times in SHCommit\_wafl routine.

## waf1.aggr.hya.eviction.fail

### Severity

ERROR

### Description

This message occurs when Flash Pool (tm) eviction cannot remove or update a read-cached block due to the corruption of its metadata.

### Corrective Action

Run waf1ron on the containing aggregate to correct the Flash Pool metadata. For more information or assistance, contact NetApp technical support.

### Syslog Message

Eviction for aggregate %s found mismatch for SSD pvbn %llu, HDD pvbn %llu, found SSD pvbn %llu.

## Parameters

**aggr** (STRING): Aggregate name.  
**expected\_ssd\_pvbn** (LONGINT): SSD block number expected in the cache metadata for the original block.  
**hdd\_pvbn** (LONGINT): HDD block number of the original block.  
**cachemap\_ssd\_pvbn** (LONGINT): SSD block number found in the cache metadata for the original block.

## waf1.aggr.hya.full.high

### Severity

ERROR

### Description

This message occurs when a hybrid aggregate unexpectedly runs low on SSD free space. Contact NetApp technical support.

### Corrective Action

(Call support).

### Syslog Message

Hybrid aggregate %s SSD free space unexpectedly ran critically low. Contact NetApp technical support.

## Parameters

**vol** (STRING): Hybrid aggregate name.

## waf1.aggr.hya.full.low

### Severity

ERROR

## Description

This message occurs when a hybrid aggregate unexpectedly runs low on SSD free space. Contact NetApp technical support.

## Corrective Action

(Call support).

## Syslog Message

Hybrid aggregate %s SSD free space unexpectedly ran low. Contact NetApp technical support.

## Parameters

**vol** (STRING): Hybrid aggregate name.

## waf1.aggr.mcc.callbacks

### Severity

NOTICE

### Description

This message occurs during a MetroCluster(tm) aggregate switchback operation, when the system ignores outstanding oplock recalls, which can take a long time. This is to allow the switchback operation to make progress.

### Corrective Action

(None).

### Syslog Message

Ignoring recalling %d locks for aggregate %s %s%s%s%s, so that the MetroCluster switchback can make progress.

### Parameters

**count** (INT): Number of ignored locks.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the aggregate.

**aggr** (STRING): Name of the aggregate.

**app** (STRING): Application UUID.

**aggrident** (STRING): Unique identifier of the aggregate in case the aggregate name itself is insufficient.

## waf1.aggr.nospc.del.snapshot

### Severity

NOTICE

### Description

This message occurs when Snapshot copies of an aggregate are autodeleted to free up space for newer writes.

### Corrective Action

(None).

## Syslog Message

Autodeletion of Snapshot copies for aggregate "%s" was triggered due to: %s. Available space: %d%%, space used by Snapshot copies: %d%%.

## Parameters

**aggr** (STRING): Name of the aggregate.

**reason** (STRING): The reason why autodeletion of Snapshot copies was triggered.

**avail\_pct** (INT): Percentage of available space in the aggregate.

**snap\_used\_pct** (INT): Percentage of space used by Snapshot copies in the aggregate.

## wafl.aggr.overcommitted.vsm

### Severity

ERROR

### Description

This message occurs when there is not enough free space in the aggregate to guarantee free space in all the SnapMirror® destinations.

### Corrective Action

Increase the free space in the aggregate by shrinking existing FlexVol(tm) volumes or adding more disk space.

## Syslog Message

Not enough free space in aggregate %s to guarantee free space in all SnapMirror destinations.

## Parameters

**aggr\_name** (STRING): Aggregate name.

## wafl.aggr.recompact.complete

### Severity

INFORMATIONAL

### Description

This message occurs when the "aggr recompact " scanner has finished.

### Corrective Action

(None).

## Syslog Message

%s scanner finished on %s %s%s%s%s. %s tried on %ld blocks.

## Parameters

**opname** (STRING): Opname for the scanner.

**voltype** (STRING): Volume type on which scanner was run.

**aggr\_owner** (STRING): Owner of the volume.

**aggr\_name** (STRING): Volume name.

**app** (STRING): Application UUID.

**uuid** (STRING): Volume uuid .

**opname1** (STRING): Opname for the scanner.

**val** (LONGINT): Number of vbns acted on.

## **waf1.aggr.recompact.estimate**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the "aggr recompact analyze" command has finished and analyzed how many compacted blocks are worth recompacting.

### **Corrective Action**

(None).

### **Syslog Message**

Recompact analyze scan finished on aggregate %s%. Block fragmentation distribution is [%ld%%, %ld%%, %ld%%, %ld%%] for %llu compacted blocks. Average compaction ratio for the aggregate is %lu.%02lu:1 approximately.

### **Parameters**

**aggr\_owner** (STRING): Owner of the aggregate.  
**aggr\_name** (STRING): Aggregate name.  
**val\_low** (LONGINT): Percentage of blocks less than 25% full.  
**val\_mid** (LONGINT): Percentage of blocks less than 50% full.  
**val\_high** (LONGINT): Percentage of blocks less than 75% full.  
**val\_full** (LONGINT): Percentage of blocks greater than 75% full.  
**val\_total** (LONGINT): Total number of compacted blocks.  
**avg\_compaction** (LONGINT): Average compaction ratio.  
**avg\_compaction1** (LONGINT): Average compaction ratio.

## **waf1.aggr.relocation.times**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an aggregate relocation is complete to record timing information about the relocation process.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**aggrName** (STRING): Name of the aggregate.  
**timesString** (STRING): Report of aggregate relocation phase timing.

## **waf1.aggr.reserve.set**

### **Severity**

NOTICE

## Description

This message occurs during mounting a volume, boot or takeover/giveback. This message indicates that the aggregate reserve was set to the saved value.

## Corrective Action

(None).

## Syslog Message

Aggregate reserve for aggregate %s was set to %d.

## Parameters

**aggr** (STRING): Name of the aggregate.

**value** (INT): New aggregate reserve.

## wapl.aggr.reserve.short

### Severity

ERROR

### Description

This message occurs when we try to set an aggregate reserve. This message indicates that there is not enough space to honor the saved value of aggregate reserve.

### Corrective Action

Try creating space in the aggregate by shrinking or destroying volumes, adding disks or changing the guarantee type of volumes in the aggregate to NONE

### Syslog Message

Aggregate reserve for aggregate %s could not be set as we are short by %s.

### Parameters

**aggr** (STRING): Name of the aggregate.

**short\_by** (STRING): How much disk space we are short by.

## wapl.aggr.rsv.low.nomount

### Severity

EMERGENCY

### Description

This message occurs when an aggregate consumes almost all of its free space reserve and should not be brought back online without assistance from technical support to address the free space issue.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Failing mount of aggregate %s. Free space reserve is too low. Contact NetApp technical support for assistance.

## Parameters

**aggr** (STRING): Aggregate name.

## **wafI.aggr.snapReserve**

### Severity

NOTICE

### Description

This message occurs when the system boot process detects that the aggregate snapshot spill percentage has been modified by bootargs.

### Corrective Action

(None).

### Syslog Message

Aggregate snapshot spill percentage is set to %d%%.

## Parameters

**aggr\_snap\_spill\_pct** (INT): Aggregate snapshot spill percentage.

## **wafI.aggr.uncompact.complete**

### Severity

INFORMATIONAL

### Description

This message occurs when the "aggr uncompact " scanner has finished.

### Corrective Action

(None).

### Syslog Message

Uncompact scanner finished on %s %s%s%s%s. Uncompactd %ld blocks.

## Parameters

**voltype** (STRING): Volume type on which scanner was run.

**aggr\_owner** (STRING): Owner of the volume.

**aggr\_name** (STRING): Volume name.

**app** (STRING): Application UUID.

**uuid** (STRING): Volume uuid .

**val** (LONGINT): Number of vbns acted on.

## **wafI.aggr.uncompact.rep.error**

### Severity

ERROR

### Description

This message occurs when the "aggr uncompact scan" command fails to uncompact some blocks because of RAID errors. The system cannot be restored to the previous release until the problem is corrected.

## Corrective Action

Uncompaction of blocks with RAID errors can be done safely only on flex volumes. Run the "vol efficiency undo -data-compaction" command on the flex volume that reported RAID errors.

## Syslog Message

RAID errors found when uncompaction scan was run on aggregate %s%. Run uncompaction scan on flex volume %s.

## Parameters

**aggr\_owner** (STRING): Owner of the aggregate.  
**aggr\_name** (STRING): Aggregate name.  
**vvol\_name** (STRING): Flex volume name.

## wafI.aggr.uncompact.skipped

### Severity

NOTICE

### Description

This message occurs when uncompact scanner has to abort due to some error conditions.

### Corrective Action

(None).

### Syslog Message

Skipping uncompact scan on aggregate %s. Volume %s encountered error %d.

### Parameters

**aggr\_name** (STRING): Aggregate name.  
**vvol\_name** (STRING): Flex volume name.  
**error\_code** (INT): Error code for the abort.

## wafI.aggregate events

### wafI.aggregate.busy

#### Severity

NOTICE

#### Description

This message occurs when an operation is run on an aggregate while the aggregate or any volume of the aggregate is busy. The message prints the reason for aggregate or volume being busy.

#### Corrective Action

(None).

#### Syslog Message

%s: %s '%s' is busy. Reason: %s.

## Parameters

**operation** (STRING): Name of the operation.

**voltype** (STRING): Aggregate or volume.

**aggrname** (STRING): Aggregate or volume name.

**reason** (STRING): Reason for the aggregate or volume being busy.

## waf.aggregate.online.phases

### Severity

INFORMATIONAL

### Description

Issued upon successful completion of aggregate online to record the sub-phase times for different phases of online.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**aggrName** (STRING): The name of the Aggregate.

**uuid** (STRING): The UUID of the Aggregate.

**phase1** (LONGINT): Total time taken for aggregate online phase1.

**phase2** (LONGINT): Total time taken for aggregate online phase2.

**phase3** (LONGINT): Total time taken for aggregate online phase3.

## waf.aggregate.online.replay

### Severity

INFORMATIONAL

### Description

Issued upon aggr online nvlog replay completion.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**aggrName** (STRING): The name of the Aggregate.

**uuid** (STRING): The UUID of the Aggregate.

**time** (LONGINT): Total time taken for nvlog replay.

## waf.aggrtrans events



## **waf1.aggrtrans.corrupted**

### **Severity**

ALERT

### **Description**

This message occurs when an aggregate transition operation encounters a corrupted aggregate.

### **Corrective Action**

Delete the corrupted aggregate in maintenance mode. Alternatively, return the disks to a 7-Mode environment to fix the corrupted aggregate.

### **Syslog Message**

Aggregate transition encountered a corrupted aggregate %s.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **waf1.aggrtrans.dirtyFS**

### **Severity**

NOTICE

### **Description**

This message occurs when an aggregate transition operation fails because the aggregate was not cleanly shut down.

### **Corrective Action**

Retry the aggregate transition to bypass this specific error. Alternatively, return the disks to a 7-Mode environment, cleanly shut down the system, and then try the transition again.

### **Syslog Message**

Aggregate %s was not cleanly shut down and could not be transitioned.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **waf1.aggrtrans.dirtyFS.cont**

### **Severity**

NOTICE

### **Description**

This message occurs when an aggregate transition continues to operate on an aggregate that was not cleanly shut down.

### **Corrective Action**

(None).

### **Syslog Message**

Aggregate %s was not cleanly shut down and is continuing transition.

**Parameters**

**aggrname** (STRING): Name of the aggregate.

**waf1.aggrtrans.outofspace.nodereboot****Severity**

ALERT

**Description**

This message occurs when a transitioned aggregate reaches a space limitation during the precommit phase and cannot be taken offline. The node is automatically rebooted to ensure that data integrity is preserved. After the reboot, the aggregate is online and its FlexVol® volumes are offline.

**Corrective Action**

Add disks to the aggregate to ensure that enough free space is available. Alternatively, commit the current transition session.

**Syslog Message**

Transitioned aggregate %s is critically low on available space, initiating node reboot.

**Parameters**

**aggrname** (STRING): Name of the aggregate.

**waf1.aggrtrans.outofspace.normal****Severity**

NOTICE

**Description**

This message occurs when a previously full transitioned aggregate contains enough free space and is back to normal operation.

**Corrective Action**

(None).

**Syslog Message**

Transitioned aggregate %s now contains enough free space and is now back to normal.

**Parameters**

**aggrname** (STRING): Name of the aggregate.

**waf1.aggrtrans.outofspace.offline****Severity**

EMERGENCY

**Description**

This message occurs when a transitioned aggregate reaches a space limitation during the precommit phase and is taken offline.

### **Corrective Action**

Add disks to the aggregate to ensure that enough free space is available. Alternatively, commit the current transition session.

### **Syslog Message**

Transitioned aggregate %s is critically low on available space, going offline.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **waf1.aggrtrans.outofspace.remount.needed**

### **Severity**

NOTICE

### **Description**

This message occurs when a previously full transitioned aggregate is online and must be remounted to bring its volumes back online.

### **Corrective Action**

Take offline the transitioned aggregate, and then bring online the aggregate again to bring its volumes back online.

### **Syslog Message**

Transitioned aggregate %s must be remounted (offline and online) to bring its volumes back online.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **waf1.aggrtrans.outofspace.warn**

### **Severity**

ERROR

### **Description**

This message occurs when a transitioned aggregate is low on free space. The message is periodically displayed before the aggregate goes offline due to being out of space.

### **Corrective Action**

Add disks to the aggregate to resume normal operations. Alternatively, commit the current transition session.

### **Syslog Message**

Pre-commit transitioned aggregate %s is low on free space at about %d percent used.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

**percent** (INT): Percentage of aggregate that is used.

## **waf1.aggrtrans.revert.to.snap**

### **Severity**

NOTICE

### **Description**

This message occurs when an aggregate transition operation fails because the aggregate is set to revert to a Snapshot copy.

### **Corrective Action**

Return the disks to a 7-Mode environment, bring the aggregate online to finish restoring to the Snapshot® copy, and then try the transition again.

### **Syslog Message**

Aggregate %s is reverting to a Snapshot copy and cannot be transitioned.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **waf1.aggrtrans.snap.delErr**

### **Severity**

ERROR

### **Description**

This message occurs when an aggregate transition commit operation fails to delete the pinned Snapshot® copy.

### **Corrective Action**

Manually delete the pinned Snapshot copy using the node scoped command "snap delete -A". If manual deletion fails, contact technical support.

### **Syslog Message**

Aggregate transition commit operation fails to delete the pinned Snapshot® copy "%s" on aggregate %s.  
Error: %s.

### **Parameters**

**name** (STRING): Name of the Snapshot copy.  
**aggrname** (STRING): Name of the aggregate.  
**err** (STRING): Error description.

## **waf1.aggrtrans.snapshot.32b**

### **Severity**

ERROR

### **Description**

This message occurs when an aggregate transition fails due to detection of a 32-bit Snapshot copy on the transition aggregate.

### Corrective Action

Retry aggregate transition using "-ignore-errors" parameter to automatically delete all the 32-bit Snapshot copies as part of aggregate transition.

### Syslog Message

Aggregate transition failed due to unsupported 32-bit snapshots on aggregate %s uuid %s.

### Parameters

**aggrname** (STRING): Name of the aggregate.

**aggruuid** (STRING): Unique identifier of the aggregate.

## wafl.air events

### wafl.air.info

#### Severity

ERROR

#### Description

Generic wafl air message

#### Corrective Action

(None).

#### Syslog Message

%s

#### Parameters

**msg** (STRING): The message from AIR module

## wafl.analytics events

### wafl.analytics.air.uh

#### Severity

ALERT

#### Description

This message occurs when the system detects an inconsistency in the file system analytics database. As a consequence, file system analytics has been disabled.

#### Corrective Action

Turn file system analytics on again for the volume to initiate new consistent analytics.

#### Syslog Message

The system detected an inconsistency in the file system analytics database associated with volume %s:%s.

#### Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

## **waf.analytics.corruption**

### **Severity**

ALERT

### **Description**

This message occurs when a file system analytics record is corrupted. Instances of this message indicate low accuracy of the analytics data reported for the volume.

### **Corrective Action**

Please contact NetApp technical support for assistance.

### **Syslog Message**

A file system analytics record identified by analytics filehandle "%s" on the volume with DSID "%lu" was corrupted due to an unexpected error.

### **Parameters**

**afh** (STRING): Key identifies the corrupted file system analytics record.

**dsid** (LONGINT): DSID of the volume affected by the corrupted record.

## **waf.analytics.discardNotOk**

### **Severity**

ALERT

### **Description**

This message occurs when the analytics subsystem discards a pending reminder work request due to an internal error. Repeated instances of this message indicate low accuracy of the analytics data reported for the volume.

### **Corrective Action**

Taking the volume offline and bringing it back online might resolve the issue. Find the name of the Vserver and volume using the "volume show -dsid <dsid>" command. Take the volume offline using the "volume offline" command, and then bring it back online by using the "volume online" command. If the problem persists or further assistance is needed, contact NetApp technical support.

### **Syslog Message**

An analytics reminder work request was discarded due to an unexpected error. The corresponding volume dsid is "%lu". The error returned was: %s

### **Parameters**

**dsid** (LONGINT): Dsid of the volume affected by the discarded reminder work request.

**error** (STRING): Error returned by reminder processing.

## **waf.analytics.enterOverload**

### **Severity**

NOTICE

### **Description**

This message occurs when the analytics subsystem is not keeping up with the amount of work being generated by the system. This could be an indication of an overloaded system in general. Analytics

accounting might lag, and reported data might not be as current as the data that the subsystem usually reports.

#### **Corrective Action**

The system should attempt to resolve the problems by itself. If problems persist or further assistance is needed, contact NetApp technical support.

#### **Syslog Message**

The analytics subsystem is not keeping up with the amount of work being generated by the system.

#### **Parameters**

(None).

### **waf.analytics.exitOverload**

#### **Severity**

NOTICE

#### **Description**

This message occurs when the analytics subsystem is performing normally and has caught up with its pending work.

#### **Corrective Action**

(None).

#### **Syslog Message**

The analytics subsystem is performing normally.

#### **Parameters**

(None).

### **waf.analytics.init.fail**

#### **Severity**

ALERT

#### **Description**

This message occurs when a volume fails to have a file system analytics database initialized.

#### **Corrective Action**

Disable then re-enable file system analytics for the volume. If problems persist or further assistance is needed, contact NetApp technical support.

#### **Syslog Message**

Failed to initialize a file system analytics database for volume %s:%s, %s.

#### **Parameters**

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

**error** (STRING): Error message.

## waf.analytics.languishing

### Severity

ALERT

### Description

This message occurs when the analytics subsystem is unable to make progress, which could be an indication of an unhealthy node or a larger systemic issue.

### Corrective Action

NetApp technical support

### Syslog Message

The analytics subsystem has languished for more than 900 seconds.

### Parameters

(None).

## waf.analytics.offErr

### Severity

ALERT

### Description

This message occurs when file system analytics fails to turn off. This can cause incorrect file system analytics results.

### Corrective Action

Use the "volume show -dsid <dsid>" command to find the volume on which file system analytics failed to turn off, and then use the "volume analytics off" command to attempt to turn off file system analytics. If problems persist or further assistance is needed, contact NetApp technical support.

### Syslog Message

File system analytics failed to turn off on the volume with DSID "%lu". Error: %d. Context: %s.

### Parameters

**dsid** (LONGINT): Data Set ID (DSID) of the volume on which file system analytics failed to turn off.

**error** (INT): Error code returned when file system analytics failed to turn off.

**context** (STRING): Which operation attempted to turn off the file system analytics.

## waf.analytics.recover

### Severity

ALERT

### Description

This message occurs when a volume fails to recover the File System Analytics state after failing to start the initialization scanner. As a consequence, File System Analytics information might be unavailable or might not accurately reflect the volume's content.



### Corrective Action

Use the "volume analytics off" command to disable File System Analytics on the volume, and then use the "volume analytics on" command to enable File System Analytics. If problems persist or further assistance is needed, contact NetApp technical support.

### Syslog Message

The File System Analytics state for volume DSID (%lu) was not recovered after the initialization scanner failed to start. Reason: %s.

### Parameters

**dsid** (LONGINT): Unique identifier (DSID) of the volume.  
**error** (STRING): Error message.

## waf.analytics.rescan

### Severity

ALERT

### Description

This message occurs when the system detects an inconsistency in the file system analytics database. As a consequence, file system analytics information might be unavailable or might not accurately reflect the volume's content.

### Corrective Action

Turn file system analytics off and then on again for the volume. If problems persist or further assistance is needed, contact NetApp technical support.

### Syslog Message

The system detected an inconsistency in the file system analytics database associated with volume %s:%s.

### Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.  
**volname** (STRING): Volume name.

## waf.analytics.scan.fail

### Severity

ALERT

### Description

This message occurs when the "volume analytics on" command fails due to an internal error.

### Corrective Action

Resolve the error if possible. Disable file system analytics using the "volume analytics off" command and then re-enable file system analytics using the "volume analytics on" command. If problems persist or further assistance is needed, contact NetApp technical support.

### Syslog Message

File system analytics scan failed for volume "%s:%s. Error: %s".

## Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

**error** (STRING): Error message.

## wafl.analytics.state.change

### Severity

NOTICE

### Description

This message occurs when the file system analytics state has been changed on a volume.

### Corrective Action

(None).

### Syslog Message

The file system analytics state of volume "%s:%s" (DSID: "%lu") has changed to "%s". Reason: %s.  
Rulegen: %lu.

## Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's storage VM (SVM).

**volname** (STRING): Volume name.

**dsid** (LONGINT): Unique identifier (DSID) of the volume.

**state** (STRING): File system analytics state: on or off.

**reason** (STRING): The reason that the file system analytics state was changed.

**rulegen** (LONGINT): File system analytics rules generation number.

## wafl.analytics.tooMuchMemory

### Severity

ALERT

### Description

This message occurs when the analytics subsystem has an excessive amount of pending work, causing memory usage to be higher than expected.

### Corrective Action

Contact Contact NetApp technical support. for assistance.

### Syslog Message

Reminder table size is "%lu", which is above the design limit of "%lu". System memory: "%lu". Time since last merge: "%lu".

## Parameters

**tablesize** (LONGINT): Current size of analytics subsystem reminder table.

**limit** (LONGINT): Design limit size of analytics subsystem reminder table.

**system\_memory** (LONGINT): Total system memory on this node.

**last\_merge** (LONGINT): Milliseconds since the analytics subsystem last processed a work request.

## wafl.analytics.tz.changed

### Severity

ERROR

### Description

This message occurs when the system timezone is changed while file system analytics is enabled on the node. As a consequence, file system analytics information might not accurately reflect the volume's content.

### Corrective Action

Use the "volume analytics show -state on" command to find the volumes with file system analytics enabled, then use the "volume analytics off" and "volume analytics on" commands for each volume where file system analytics is enabled.

### Syslog Message

The system timezone has been changed while file system analytics is enabled.

### Parameters

(None).

## wafl.analytics.updateDiscard

### Severity

ALERT

### Description

This message occurs when a file system analytics update is discarded. Instances of this message indicate low accuracy of the analytics data reported for the volume.

### Corrective Action

Turning file system analytics off and turning it back on again might resolve the issue. Find the name of the Vserver and volume using the "volume show -dsid <dsid>" command. Turn file system analytics off using the "volume analytics off" command, and then turn it back on by using the "volume analytics on" command. If the problem persists or further assistance is needed, contact NetApp technical support.

### Syslog Message

A file system analytics update was discarded due to an unexpected error. Analytics filehandle "%s" on the volume with DSID "%lu" was to be updated by file operation "%s".

### Parameters

**fh** (STRING): File handle that the operation targeted.

**dsid** (LONGINT): DSID of the volume affected by the discarded update.

**operation** (STRING): Operation that was being performed when this error occurred.

## wafl.ap events

### wafl.ap.info

### Severity

NOTICE

**Description**

adaptive playlist message

**Corrective Action**

(None).

**Syslog Message**

%s

**Parameters**

**msg** (STRING): The message from wafI adaptive playlist

## wafI.arI events

### wafI.arI.vol.limitChkFailed

**Severity**

ERROR

**Description**

This message occurs during the relocation of an aggregate, when the destination node fails the volume count limit and volume size limit checks.

**Corrective Action**

Ensure that the destination node supports the volume size and volume count limits, or select a different destination node and retry the operation.

**Syslog Message**

Aggregate '%s': Volume limits check failed on destination node %s.

**Parameters**

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## wafI.asserts events

### wafI.asserts.disabled

**Severity**

INFORMATIONAL

**Description**

This message occurs on development kernels when some of the asserts are disabled.

**Corrective Action**

(None).

**Syslog Message**

Warning: Some of the asserts are disabled on the filer.

## Parameters

(None).

# wabl.atime events

## wabl.atime.updates.restarted

### Severity

NOTICE

### Description

This message occurs when the access time updates of all the files in the volumes of an aggregate start because the aggregate has enough free space. Access time updates of all the files were stopped when the aggregate ran out of space. The aggregate now has enough space to start the access time updates.

### Corrective Action

(None).

### Syslog Message

Access time updates of all the files in the volumes of aggregate "%s" have started.

### Parameters

**aggregate\_name** (STRING): Aggregate name for which access time updates of all the files in the volumes start.

## wabl.atime.updates.stopped

### Severity

ALERT

### Description

This message occurs when the access time updates of all the files in the volumes of an aggregate stop because the aggregate has consumed its entire free space.

### Corrective Action

Create space by increasing the aggregate size or by deleting some volumes or moving volumes to other aggregates. Use the "storage aggregate show-space" command to view the current space usage of the aggregate. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot® copies are deleted automatically to make some space when the aggregate is full.

### Syslog Message

Access time updates of all the files in the volumes of aggregate "%s" have been stopped.

### Parameters

**aggregate\_name** (STRING): Aggregate name for which access time updates of all the files in the volumes stop.

## wabl.atime.updates.will.stop

**Severity**

ALERT

**Description**

This message occurs when the aggregate is nearly full. The access time updates of all the files in the volumes of an aggregate stop when the aggregate consumes its entire free space.

**Corrective Action**

Create space by increasing the aggregate size or by deleting some volumes or moving volumes to other aggregates. Use the "storage aggregate show-space" command to view the current space usage of the aggregate. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot® copies are deleted automatically to make some space when the aggregate is full.

**Syslog Message**

Access time updates of all the files in the volumes of aggregate "%s" stop when the aggregate becomes full.

**Parameters**

**aggregate\_name** (STRING): Aggregate name for which access time updates of all the files in the volumes stop.

## waf.attr events

### waf.attr.topK.clear.fail

**Severity**

ALERT

**Description**

This message occurs when the large directory report fails to be cleared for a volume when this feature was turned off due to an internal error. Report on large directories may be unavailable or incorrect when it is turned on again.

**Corrective Action**

The system should attempt to resolve the problems by itself. If problems persist or further assistance is needed, contact NetApp technical support.

**Syslog Message**

Clearing the large directory report failed for volume "%s:%s. Error: %s"

**Parameters**

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

**error** (STRING): Error message.

### waf.attr.topK.init.fail

**Severity**

ALERT

## Description

This message occurs when the large directory report fails to be initialized for a volume due to an internal error. Report on large directories may be unavailable or incorrect.

## Corrective Action

The system should attempt to resolve the problems by itself. If problems persist or further assistance is needed, contact NetApp technical support.

## Syslog Message

Initializing the large directory report failed for volume "%s:%s. Error: %s"

## Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

**error** (STRING): Error message.

## wafl.attr.topK.scan.fail

### Severity

ALERT

### Description

This message occurs when scanning analytics database for large directories data fails due to an internal error. Report on large directories may be unavailable or incorrect.

### Corrective Action

The system should attempt to resolve the problems by itself. If problems persist or further assistance is needed, contact NetApp technical support.

### Syslog Message

Scanning analytics database for large directory report failed for volume "%s:%s. Error: %s"

### Parameters

**vserveruuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volname** (STRING): Volume name.

**error** (STRING): Error message.

## wafl.bad events

### wafl.bad.aggr.buftree.type

#### Severity

EMERGENCY

#### Description

This message occurs when you attempt to bring an aggregate online whose file system buftree type and RAID label format are inconsistent. Buftree type indicates whether the filesystem is a 32-bit or a 64-bit. Attempting to online this aggregate can cause data corruption or a system panic.

## Corrective Action

Contact NetApp technical support

## Syslog Message

Aggregate %s: Inconsistency found with the file system buftree type and RAID label format (large aggr: %d, fs: 0x%x, RAID: %d).

## Parameters

**name** (STRING): The aggregate name.  
**large** (INT): Flag showing if the aggregate size is bigger than 16 TB.  
**fsbuftree** (INT): File system buftree type.  
**raidfmt** (INT): RAID buftree format.

## waf1.bad.vol.buftree.type

### Severity

EMERGENCY

### Description

This message occurs when you attempt to bring online a volume whose file system and volume descriptor buftree type show inconsistency. The buftree type shows whether the file system is a 32-bit or 64-bit. Bringing online such a volume can cause data corruption or system panic.

### Corrective Action

(Call support).

## Syslog Message

Volume %s: Inconsistency found with the file system buftree type and volume descriptor buftree type (large vol: %d, fs: 0x%x, label: 0x%x, RAID: %d). Bringing the volume online can cause data corruptions or system panic. Contact NetApp technical support.

## Parameters

**volname** (STRING): Volume name.  
**large** (INT): Flag showing if the volume size is bigger than 16 TB.  
**fsbuftree** (INT): File system buftree type.  
**labelfmt** (INT): Volume descriptor buftree type.  
**raidfmt** (INT): RAID buftree type.

## waf1.bagofbits events

### waf1.bagofbits.name

### Severity

ERROR

### Description

This message occurs when a read directory request from an NFSv4 client is made to a Unicode-based directory in which directory entries with no NFS alternate name contain non-Unicode characters.



## Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Access the parent directory from an NFSv3 client and rename the entry using Unicode characters.

## Syslog Message

Cannot provide directory entry to NFSv4 client for file ID:%d in volume %s%s because it contains non-Unicode characters with no NFS alternate name.

## Parameters

**fileid** (INT): File ID of the entry in the directory for which non-Unicode characters are present.  
**vol** (STRING): Volume in which the directory entry with non-unicode characters is present.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identifier (UUID) when the volume name itself is insufficient.

# wافل.blkreclaim events

## wافل.blkreclaim.abort

### Severity

INFORMATIONAL

### Description

This message occurs to inform you that the block reclamation scanner was aborted.

### Corrective Action

You can disable this message as follows: `setflag wافل_blk_reclaim_scan_log 0`

### Syslog Message

Block reclamation scanner was aborted on %s %s%s%s%s.

### Parameters

**type** (STRING): Type of object (volume/aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume when the volume name alone is insufficient.

## wافل.blkreclaim.complete

### Severity

INFORMATIONAL

### Description

This message occurs when the block reclamation scanner is complete.

### Corrective Action

Disable this message by running the '`setflag wافل_blk_reclaim_scan_log 0`' command.

### Syslog Message

Block reclamation scanner was completed on %s %s%s%s%s. The scanner took %ld ms.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.  
**run\_time** (LONGINT): Total run time of the scanner.

## waf1.blkreclaim.noscan

### Severity

INFORMATIONAL

### Description

This event is issued to inform the user that the block reclamation scanner was not started.

### Corrective Action

This message may be disabled with: `setflag waf1_blk_reclaim_scan_log 0`

### Syslog Message

Block reclamation scanner not started on %s %s%s%s%s.

## Parameters

**type** (STRING): The type of object (volume/aggregate)  
**owner** (STRING): The volume owner  
**vol** (STRING): The volume name  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.blkreclaim.reset

### Severity

INFORMATIONAL

### Description

This message occurs to inform you that the block reclamation scanner was reset.

### Corrective Action

You can disable this message as follows: `setflag waf1_blk_reclaim_scan_log 0`

### Syslog Message

Block reclamation scanner was reset on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume/aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume when the volume name alone is insufficient.

# waf1.br events

## waf1.br.redir.invalid

### Severity

ALERT

### Description

This event is issued if block reallocation has not been performed on an aggregate but a reallocated block was found that subsequently needs to be redirected. Block reallocation operates in two phases: first the blocks on an aggregate are moved (reallocated), later, when the block is read, the old pointers to the block are updated (redirected). This event is issued when the second phase is required without the first phase occurring. Since this may indicate on-disk data corruption, the aggregate has been taken offline.

### Corrective Action

(Call support).

### Syslog Message

Invalid redirection: aggr:%s vol:%s

### Parameters

**aggr** (STRING): The name of the aggregate

**vol** (STRING): The name of the volume

## waf1.br.redir.nospace

### Severity

NOTICE

### Description

This message occurs when there's not enough free space to complete a redirect scan on a volume. The redirect scan improves performance by updating information that was changed because of block reallocation, but is not required for normal operation. This message is usually caused because reallocated blocks are shared between the active file system and a Snapshot(tm) copy and cannot be updated without using more blocks than are currently available.

### Corrective Action

Create more free space in the volume by deleting Snapshot copies.

### Syslog Message

The redirect scanner cannot complete its scan on '%s%s%s' because there is not enough disk space.  
Remove Snapshot copies.

### Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier the volume in cases where the volume name itself is insufficient.

# waf1.buf.events

## waf1.buf.badHeader

### Severity

EMERGENCY

### Description

This message occurs when the system encounters a bad header in the Snapshot(tm) copy of the active map block of the filesystem.

### Corrective Action

Delete the bad Snapshot copy. For a SnapMirror® target, rebuild from source or resynchronize, after deleting the Snapshot copy. If the problem persists or if you need further assistance, contact NetApp technical support.

### Syslog Message

Bad header in Snapshot copy of active map in %s %s%s%s%s, Snapshot copy %d, fbn %llu.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**snapld** (INT): Snapshot(tm) copy identifier.

**fbn** (LONGINT): File block number.

## waf1.buf.freeingFreeBlock

### Severity

EMERGENCY

### Description

This message occurs when the system encounters a "freeing free block" error due to file system corruption.

### Corrective Action

(Call support).

### Syslog Message

Freeing free block at VBN %llu in %s %s%s%s%s. WAFL inconsistent. Contact NetApp technical support.

### Parameters

**vbn** (LONGINT): Virtual block number of the block that is freed.

**voltype** (STRING): Type of the volume.

**volowner** (STRING): Owner of the volume.

**volname** (STRING): Name of the volume (again). This is returned from WAFL\_VOL\_PRINT\_NAME\_STRINGS.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**vendor** (STRING): Name of the vendor.

# waf.c2c events

## waf.c2c.increase.vol.max.size.failed

### Severity

ERROR

### Description

This message occurs when attempt to increase volume max size limit during c2c restore fails due to an error.

### Corrective Action

(None).

### Syslog Message

Attempt to increase max size limit for volume:%s failed with Error:%d.

### Parameters

**vvolname** (STRING): Name of the volume.

**errorCode** (INT): Error code

## waf.c2c.od.restore.scan.cloud.config.failed

### Severity

NOTICE

### Description

This message occurs when cloud config is unavailable and SMC on-demand restore from object store failed to start.

### Corrective Action

(None).

### Syslog Message

Phase:%s of the on-demand restore process for volume:%s with FSID:0x%x has failed. Attempt to restart SMC on-demand restore on this volume will continue in the background.

### Parameters

**restore\_phase** (STRING): Specifies the phase for which the restore has been failed.

**vol** (STRING): Volume name.

**fsid** (INT): File system identifier of the volume.

## waf.c2c.od.restore.scan.initiated

### Severity

NOTICE

### Description

This message occurs when an on-demand restore from object store has been initiated.

### Corrective Action

(None).

### Syslog Message

Phase:%s of the on-demand restore process for volume:%s with FSID:0x%x has initiated.

### Parameters

**restore\_phase** (STRING): Specifies the phase for which the restore has been initiated.

**vol** (STRING): Volume name.

**fsid** (INT): File system identifier of the volume.

## waf1.c2c.od.restore.scan.restart.failed

### Severity

ALERT

### Description

This message occurs when SMC lazy restore scan restart attempt for transient error failed.

### Corrective Action

(None).

### Syslog Message

Restart Phase:%s of the on-demand restore process for volume:%s with FSID:0x%x has failed. For more information or assistance, contact NetApp technical support.

### Parameters

**restore\_phase** (STRING): Specifies the phase for which the restore has been failed.

**vol** (STRING): Volume name.

**fsid** (INT): File system identifier of the volume.

## waf1.c2c.restore.scan.abort

### Severity

NOTICE

### Description

This message occurs when an operation to restore file(s)/volume from a cloud storage was aborted.

### Corrective Action

(None).

### Syslog Message

%s %s%s%s, %s Restore from cloud storage aborted.

### Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**restore\_type** (STRING): Indicates whether the operation is Full Volume, Single File or On-demand Restore.

## **waf.c2c.restore.scan.done**

### **Severity**

NOTICE

### **Description**

This message occurs when an operation to restore file(s)/volume from a cloud storage was completed successfully.

### **Corrective Action**

(None).

### **Syslog Message**

%s %s%s%s, %s Restore from cloud storage completed.

### **Parameters**

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**restore\_type** (STRING): Indicates whether the operation is Full Volume, Single File or On-demand Restore.

## **waf.c2c.restore.scan.failed**

### **Severity**

ERROR

### **Description**

This message occurs when there is an error during restoring of file(s) or Volume from a Cloud storage.

### **Corrective Action**

(None).

### **Syslog Message**

Restoring file with inode %d from cloud has failed with error %d for the volume %s with FSID 0x%x during %s Restore.

### **Parameters**

**fileid** (INT): File identifier.

**error** (INT): Error Code.

**vol** (STRING): Volume name.

**fsid** (INT): File system identifier of the volume.

**restore\_type** (STRING): Indicates whether the operation is Full Volume, Single File or On-demand Restore.

## wafl.c2c.rst.fsize.exceeded

### Severity

NOTICE

### Description

This message occurs when the size of a file being restored exceeds the maximum file size allowed by the system during the restoration of either files or a volume from cloud storage. The restore operation will fail for that particular file or volume.

### Corrective Action

(None).

### Syslog Message

%s %s%s%s%s, FSID 0x%x, Restoring file with file ID %d from cloud has failed because the file size exceeded the maximum file size allowed on the current system, which is %lu bytes. The size of the file being restored is %lu bytes.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume if the volume name itself is insufficient.

**fsid** (INT): File system identifier of the volume.

**fileid** (INT): File identifier.

**maxsize** (LONGINT): Maximum file size allowed on this system.

**filesize** (LONGINT): Size of the file being restored.

## wafl.ca events

### wafl.ca.cap.tier.usage.limit

### Severity

ERROR

### Description

This message occurs when the space used in the capacity tier of a FabricPool aggregate has exceeded or is projected to exceed the size limit set on the capacity tier. The projected usage is calculated based on the cold data already identified for tiering. The system will continue to serve data while space exists in the performance tier, but cold data will not be tiered to the capacity tier.

### Corrective Action

Perform one or more of the following actions: o Free space in the capacity tier of the aggregate by moving volumes, removing volumes, or deleting data. o Increase the size limit set on the capacity tier.

### Syslog Message

Aggregate %s capacity tier usage (%llu MB) has exceeded its size limit (%llu MB).

### Parameters

**aggregate** (STRING): Aggregate name.



**capacityTierUsage** (LONGINT): Current capacity tier usage in MB.

**capacityTierSizeLimit** (LONGINT): Capacity tier size limit in MB.

## **waf1.ca.cloud.retrieval.policy**

### **Severity**

NOTICE

### **Description**

This message occurs when the cloud retrieval policy for a volume in a FabricPool has been modified.

### **Corrective Action**

(None).

### **Syslog Message**

Cloud retrieval policy on volume '%s%s%s%s' has been modified to %s.

### **Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Identifier of the volume.

**cloud\_retrieval\_policy** (STRING): Cloud retrieval policy configured on the volume.

## **waf1.ca.cloud.write.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the cloud write option is disabled on a volume in a FabricPool aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

The Cloud Write option is disabled on volume %s.

### **Parameters**

**vol** (STRING): Name of the volume.

## **waf1.ca.cloud.write.enabled**

### **Severity**

NOTICE

### **Description**

This message occurs when the cloud write option is enabled on a volume in a FabricPool aggregate.

### **Corrective Action**

(None).

## Syslog Message

The Cloud Write option is enabled on volume %s.

## Parameters

**vol** (STRING): Name of the volume.

## waf1.ca.cs.hdr.cfg.mismatch

### Severity

ALERT

### Description

This message occurs when the system detects that the IDs recorded in the configuration map for FabricPool object stores for a specific aggregate do not match the configuration IDs currently in use by the aggregate itself. This mismatch might affect the reporting of statistics and counters for FabricPool object stores.

### Corrective Action

Run the "aggr waf1iron start <aggr\_name>" command on the aggregate specified to fix the inconsistency.

## Syslog Message

Aggregate "%s" with bin ID "%d" has configuration IDs in the FabricPool statistics file (configuration ID 1: %d, configuration ID 2: %d) that are different from the configuration IDs of the object stores in the configuration map (primary configuration ID: %d, mirror configuration ID: %d).

## Parameters

**aggr** (STRING): Name of the aggregate.

**bin** (INT): Bin ID of the object stores.

**cfg\_id\_1** (INT): Configuration ID 1, as reported by the aggregate.

**cfg\_id\_2** (INT): Configuration ID 2, as reported by the aggregate.

**primary\_cfg\_id** (INT): Configuration ID of the primary object store.

**mirror\_cfg\_id** (INT): Non-zero configuration ID of the mirror object store, or zero if no such store exists.

## waf1.ca.latency.threshold

### Severity

NOTICE

### Description

This message occurs when the measured I/O latency from a cloud provider, for a specific aggregate and object store, exceeds the maximum threshold. This can impact performance and lead to request timeouts.

### Corrective Action

(None).

## Syslog Message

Measured latency from cloud (%d msec) more than threshold (%d msec) (aggregate: "%s", aggregate uuid: "%s", object store: "%s")

## Parameters

**latency\_msec** (INT): Latency cycles specified in milliseconds.

**threshold\_msec** (INT): Threshold value in milliseconds for the latency.

**aggr** (STRING): Name of the containing aggregate.  
**aggr\_uuid** (STRING): UUID of the aggregate.  
**config\_name** (STRING): Object store configuration name.

## **waf1.ca.mcc.tlog.drain.event**

### **Severity**

NOTICE

### **Description**

This message occurs when tiering is enabled or disabled on a FabricPool in a MetroCluster(tm) configuration, when one or more mirrored plexes transitions from normal to degraded state or vice versa. Enabling or disabling tiering does not impact data availability or client read/write performance.

### **Corrective Action**

(None).

### **Syslog Message**

Mirrored plex "%s" on FabricPool "%s" is in "%s" state. Tiering is %s.

### **Parameters**

**plex\_name** (STRING): Name of the plex.  
**aggregate\_name** (STRING): Name of the FabricPool.  
**plex\_state** (STRING): Operational state of the plex.  
**message** (STRING): Indicates whether tiering is enabled or disabled.

## **waf1.ca.migrate.aggr.nospace**

### **Severity**

NOTICE

### **Description**

This message occurs when one or more FabricPool aggregates does not have space for hot data to be migrated back to the performance tier.

### **Corrective Action**

Either free space in the performance tier of the aggregate (by moving volumes, removing volumes, or deleting data), or add space to the performance tier of the aggregate by adding disks.

### **Syslog Message**

Aggregate %s: Current available space in the performance tier is below %d%%. Data migration from the capacity tier to the performance tier is stopped.

### **Parameters**

**aggr** (STRING): Aggregate name.  
**threshold\_percentage** (INT): Percentage of aggregate space that must be unused for migration to occur.

## **waf1.ca.name.fmt.mismatch**

### **Severity**

ERROR

## Description

This message occurs when the system detects an internal discrepancy on the name format version used by the FabricPool. Cloud operations on the FabricPool is affected.

## Corrective Action

Run the 'debug vreport show' command to see the object name format version discrepancy. Follow the vreport guidance to fix the issue.

## Syslog Message

Name format version discrepancy detected for FabricPool "%s". Operations to the object store affected.

## Parameters

**aggr** (STRING): Name of the FabricPool.

## waf1.ca.object.check.done

### Severity

INFORMATIONAL

### Description

This event is issued to indicate that the object check operation on a FabricPool aggregate completed successfully.

### Corrective Action

(None).

### Syslog Message

Object check operation completed on FabricPool %s. Total lost objects: %d.

### Parameters

**aggr** (STRING): Name of the containing aggregate.  
**num\_lost\_objs** (INT): Number of lost objects identified.

## waf1.ca.object.check.start

### Severity

INFORMATIONAL

### Description

This message occurs when object check operation on a FabricPool aggregate is started successfully.

### Corrective Action

(None).

### Syslog Message

Object check operation was started on FabricPool %s. Output will be saved at %s.

### Parameters

**aggr** (STRING): Name of the containing aggregate.  
**logfile** (STRING): Path for the logfile for object check output.

## **waf1.ca.object.storage.class.updated**

### **Severity**

NOTICE

### **Description**

This message occurs when the object storage class is updated on a volume in a FabricPool aggregate.

### **Corrective Action**

(None).

### **Syslog Message**

The tiering storage class on volume %s is updated from %s to %s.

### **Parameters**

**vol** (STRING): Name of the volume.

**old\_storage\_class** (STRING): Old tiering storage class.

**new\_storage\_class** (STRING): Updated tiering storage class.

## **waf1.ca.objtag.retry.limit**

### **Severity**

NOTICE

### **Description**

This message occurs when the object tagging scanner fails to update tags on an object in the object store even after retrying for the maximum allowed attempts. The scanner is aborted for the current scan and will attempt to tag the object again the next time the scanner runs. The scanner is invoked every minute to check whether there are objects to be tagged.

### **Corrective Action**

(None).

### **Syslog Message**

Failed to update object tags in the object store even after multiple retry attempts, aggregate = %s, object ID =%lu, object name =%s, retries attempted = %d.

### **Parameters**

**aggregate\_name** (STRING): Aggregate name.

**object\_id** (LONGINT): Object ID.

**object\_name** (STRING): Object name.

**num\_retries** (INT): Number of retries attempted.

## **waf1.ca.objtag.skip.retry**

### **Severity**

NOTICE

### **Description**

This message occurs when the object tagging scanner fails to update tags on an object in the object store even after retrying for the maximum allowed attempts. The scanner continues to update tags on other objects instead of aborting the scan due to repeated failures over one or more objects.

## Corrective Action

(None).

## Syslog Message

Failed to update object tags in the object store even after multiple retry attempts. Tagging skipped on object: aggregate = %s, object ID =%lu, object name =%s, retries attempted = %d.

## Parameters

**aggregate\_name** (STRING): Aggregate name.  
**object\_id** (LONGINT): Object ID.  
**object\_name** (STRING): Object name.  
**num\_retries** (INT): Number of retries attempted.

## waf1.ca.objtag.unsupported

### Severity

NOTICE

### Description

This message occurs when object tags are added or updated on a FabricPool volume. This message indicates that objects tiered from this volume will not be tagged in the object store.

### Corrective Action

(None).

### Syslog Message

Object tagging is supported only for a StorageGRID® object store in this version of ONTAP® software. Objects in %s object store (%s) attached to FabricPool %s will not be tagged in this version.

### Parameters

**provider\_name** (STRING): Object store name.  
**object\_store\_type** (STRING): Object store type (primary or mirror).  
**aggregate** (STRING): Name of the FabricPool.

## waf1.ca.promote.none.nospc

### Severity

NOTICE

### Description

This message occurs when a FabricPool volume with tiering policy set to "none" and cloud retrieve policy set to "promote" does not have enough free space available to fetch all the data from the cloud. The scanner to fetch the data will abort and be periodically rescheduled until the space is available.

### Corrective Action

Ensure the aggregate has the estimated amount of space needed available.

### Syslog Message

Aggregate %s, Volume %s: Current available space in the performance tier is %llu bytes. There needs to be at least %llu bytes free in the aggregate to fetch an estimated %llu bytes from the cloud. Data migration from the capacity tier to the performance tier is paused.

## Parameters

**aggr** (STRING): Aggregate name.

**vol** (STRING): Volume name.

**avail\_space** (LONGINT): Free space in aggregate.

**space\_needed** (LONGINT): Estimated amount of physical space that must be free to accommodate data fetched from the cloud.

**cloud\_space\_to\_fetch** (LONGINT): Estimated amount of physical space to fetch from the cloud.

## waf1.ca.promote.snap.nospc

### Severity

NOTICE

### Description

This message occurs when a FabricPool volume with tiering policy set to "snapshot\_only" and cloud retrieve policy set to "promote" does not have enough free space available to fetch the AFS data from the cloud. The scanner to fetch the data will abort and be periodically rescheduled until the space is available.

### Corrective Action

Ensure the aggregate has the estimated amount of space needed available.

### Syslog Message

Aggregate %s, Volume %s: Current available space in the performance tier is %llu bytes. There needs to be at least %llu bytes free in the aggregate to fetch AFS data from the cloud. Data migration from the capacity tier to the performance tier is paused.

## Parameters

**aggr** (STRING): Aggregate name.

**vol** (STRING): Volume name.

**avail\_space** (LONGINT): Free space in aggregate.

**space\_needed** (LONGINT): Minimum amount of physical space that must be free to accommodate AFS data fetched from the cloud.

## waf1.ca.remove.object.tag

### Severity

NOTICE

### Description

This message occurs when the system encounters a bad header in the block that stores the tag values associated with the objects that belong to the volume. As a result of this corruption, all the tags information stored in the block is cleared.

### Corrective Action

(None).

### Syslog Message

Clearing all tags due to inconsistencies found in the object tag metafile associated with volume %s.

## Parameters

**volname** (STRING): Volume name.

## waf1.ca.repaired.metafile

### Severity

ERROR

### Description

This message occurs when an inconsistency is discovered in one of the FabricPool metafile blocks. The inconsistency is automatically repaired upon detection.

### Corrective Action

(None).

### Syslog Message

Inconsistent block (%s) found in %s metafile on FabricPool aggregate (%s). Repaired.

### Parameters

**state** (STRING): A description of the metafile block that was found to be inconsistent.

**metafile\_name** (STRING): Name of the metafile.

**aggr** (STRING): Name of the FabricPool aggregate.

## waf1.ca.resync.complete

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® completes the resync process from the primary object store to the mirror object store for a mirrored FabricPool aggregate.

### Corrective Action

(None).

### Syslog Message

FabricPool mirror resync process is completed for FabricPool %s (uuid %s) from primary object store (config id %u) to mirror object store (config id %u).

### Parameters

**aggr** (STRING): Name of the containing aggregate.

**aggr\_uuid** (STRING): Universal Unique ID (UUID) of the aggregate.

**primary\_config\_id** (INT): Primary object store config id.

**mirror\_config\_id** (INT): Mirror object store config id.

## waf1.ca.resync.read.fail

### Severity

ALERT

### Description

This message occurs when FabricPool mirror resynchronize attempt receives too many consecutive timeout failures while reading a single object from the primary. The mirror will stay in a degraded state until all objects can be read and copied to the mirror.



## Corrective Action

Verify that your intercluster LIF is online and functional by using the "network interface show" command. Also, check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. Additionally, ensure that the configuration of your object store has not changed, and that login and connectivity information is still accurate. If the condition persists, contact NetApp technical support.

## Syslog Message

Unable to read cloud object from capacity tier (objid %llu, btuuid %s, seqno %llu) during FabricPool mirror resync.

## Parameters

**object\_id** (LONGINT): Object identifier.

**btuuid** (STRING): Buftree UUID.

**seqno** (LONGINT): Sequence number.

## waf1.ca.resync.read.set.err.obj

### Severity

ALERT

### Description

This message occurs when Data ONTAP® marks all blocks in an object in a capacity tier object with an error due to FabricPool mirror resync read errors. This is useful for tracking generation of some pseudobad blocks due to resync read issues.

### Corrective Action

Contact Contact NetApp technical support..

## Syslog Message

Unable to read cloud object from capacity tier (objid %llu, btuuid %s, seqno %llu) during FabricPool mirror resync. As a result, all blocks in this object are marked with error (%u).

## Parameters

**object\_id** (LONGINT): Object identifier.

**btuuid** (STRING): Buftree UUID.

**seqno** (LONGINT): Sequence number.

**error** (INT): Error on the block.

## waf1.ca.revert.ufmt.paused

### Severity

ERROR

### Description

This message occurs when the revert operation to convert a unified object format is paused either because there are one or more offline volumes in the FabricPool aggregate or because the tiering of unified format objects to the object store is pending.

### Corrective Action

Bring the volumes online in the aggregate to resume the unified object format revert operation. Use the "volume object-store tiering show" command to check the tiering status of a volume.

## Syslog Message

The operation to revert the object format to the native format paused on aggregate "%s" because there are offline volumes or the tiering of unified format objects to the object store is pending.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

## waf1.ca.slothdr.repair.done

### Severity

INFORMATIONAL

### Description

This event is issued to indicate that the slothdr repair operation for a FabricPool aggregate completed successfully.

### Corrective Action

(None).

## Syslog Message

Slothdr metafile repair operation completed for FabricPool %s.

## Parameters

**aggr** (STRING): Name of the containing aggregate.

## waf1.ca.slothdr.repair.start

### Severity

INFORMATIONAL

### Description

This message occurs when an operation to repair the slothdr metafile for a FabricPool aggregate is started successfully.

### Corrective Action

(None).

## Syslog Message

Slothdr metafile repair operation was started for FabricPool %s.

## Parameters

**aggr** (STRING): Name of the containing aggregate.

## waf1.ca.tier.aggr.spc.limit

### Severity

ERROR

### Description

This message occurs when the FabricPool aggregate does not have enough free space to tier cold data to the capacity tier. The system will continue to serve data, but cold data will not be tiered.

## Corrective Action

Either free space in the performance tier of the aggregate by moving volumes, removing volumes, or deleting data, or add space to the performance tier of the aggregate by adding disks.

## Syslog Message

Aggregate %s: Current available space in the performance tier is below %d%%. Data tiering from the performance tier is stopped.

## Parameters

**aggr** (STRING): Aggregate name.

**threshold\_percentage** (INT): Percentage of aggregate space that must be unused for tiering to occur.

## wapl.ca.tier.max.size.limit

### Severity

ERROR

### Description

This message occurs when the capacity tier of a FabricPool aggregate reaches the maximum supported capacity tier size. The system will continue to serve data, but cold data will not be migrated to the capacity tier.

### Corrective Action

This state can be corrected by performing one or more of the following actions: o Free space in the capacity tier of the aggregate by moving volumes, removing volumes, or deleting data.

### Syslog Message

Aggregate %s capacity tier usage (%llu MB) has reached the maximum capacity tier size limit (%llu MB).

### Parameters

**aggregate** (STRING): Aggregate name.

**capacityTierUsage** (LONGINT): Current capacity tier usage in MB.

**maxCapacityTierSizeLimit** (LONGINT): Maximum capacity tier size in MB.

## wapl.ca.tiering.policy

### Severity

NOTICE

### Description

This message occurs when the tiering policy for a volume in a FabricPool has been modified.

### Corrective Action

(None).

### Syslog Message

Tiering policy on volume '%s%s%s%s%' has been modified to %s.

### Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**tiering\_policy** (STRING): Tiering policy configured on the volume.

## **waf1.ca.tlogput.modified**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the FabricPool Put operation settings are modified by an administrator.

### **Corrective Action**

(None).

### **Syslog Message**

FabricPool Put operation settings were modified. New Values, bytes-limit: %llu, latency-limit: %llu, preserved: %u,%u

### **Parameters**

**bytes\_limit** (LONGINT): The put-rate-bytes-limit value.

**latency\_limit** (LONGINT): The put-latency-limit value.

**preserve** (INT): The preserve setting value.

**can\_preserve** (INT): Non-Zero if the node supports preserving settings.

## **waf1.ca.total.space.incons**

### **Severity**

ERROR

### **Description**

This message occurs when the counter tracking the total amount of space used in an object store bucket becomes inconsistent due to other inconsistencies in the object metadata.

### **Corrective Action**

For more information or assistance, contact NetApp technical support.

### **Syslog Message**

Inconsistency was found in the total object store space used by FabricPool aggregate "%s". Skipped space adjustments have caused overcounting by as much as %llu bytes.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

**space\_delta** (LONGINT): The estimated amount of space in bytes that could be overcounted as a result of other inconsistencies.

## **waf1.catalog events**

### **waf1.catalog.invalid**

**Severity**

NOTICE

**Description**

This event indicates that WAFL catalog block in nvram is invalid during nvram replay.

**Corrective Action**

(None).

**Syslog Message**

WafL Catalog invalid.

**Parameters**

(None).

## wafL.cbin events

### wafL.cbin.iron.mirror.config.changed

**Severity**

INFORMATIONAL

**Description**

This message occurs when cloud mirror configuration changes during the WAFLiron operation on a FabricPool aggregate with unified format enabled.

**Corrective Action**

Run wafLiron again on the aggregate.

**Syslog Message**

The wafLiron cannot update total physical size on bin label as cloud mirror configuration has changed on aggregate %s. Please run wafLiron again on the aggregate.

**Parameters**

**aggregate\_name** (STRING): Aggregate name.

### wafL.cbin.iron.objslothr.invalid.entry.threshold.reached

**Severity**

INFORMATIONAL

**Description**

This message occurs when number of invalid entries in objslothr metafile reaches the upper threshold during the WAFLiron operation on a FabricPool aggregate with unified format enabled.

**Corrective Action**

Run the fixup tool on the Fabricpool aggregate and after the fixup tool finishes, run wafLiron again on the aggregate.

## Syslog Message

The wafiron cannot update total physical size on bin label as number of invalid entries in slothdr metafile reached the upper threshold on aggregate %s. Please run the fixup tool on the aggregate before running wafiron again.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

## wafl.cbin.label.total.physical.size.overcounted

### Severity

INFORMATIONAL

### Description

This message occurs when overcounting of total physical size on the bin label happens due to invalid entries in slothdr metafile belonging to a Fabricpool aggregate with unified format enabled.

### Corrective Action

Run the fixup tool on the Fabricpool aggregate and after the fixup tool finishes, run wafiron again on the aggregate.

## Syslog Message

The total physical size on bin label belonging to Fabricpool aggregate %s might be overcounted as there were %llu invalid entries in slothdr metafile. The current physical used is %llu. Previously physical used was %llu. For further assistance, contact NetApp technical support.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

**invalid\_entry\_count** (LONGINT): Number of invalid entries present in slothdr metafile.

**physical\_used** (LONGINT): Current total physical used in capacity tier.

**old\_physical\_used** (LONGINT): Previous total physical used in capacity tier.

## wafl.cbin.objslothdr.iron.skip.resync.in.progress

### Severity

INFORMATIONAL

### Description

This message occurs when wafiron skips ironing objslothdr metafile on a FabricPool aggregate with unified format enabled as resync operation is in progress.

### Corrective Action

Run wafiron again once resync operation is complete.

## Syslog Message

Wafiron has skipped ironing objslothdr metafile on aggregate %s as resync operation is in progress on this aggregate. Please run wafiron after resync operation is complete.

## Parameters

**aggregate\_name** (STRING): Aggregate name.

## wافل.cde events

### wافل.cde.decompression.fail

#### Severity

ERROR

#### Description

This message occurs when the WAFL® kernel cannot decompress a compressed user data block. This is an abnormal condition that can lead to inconsistencies in the user data.

#### Corrective Action

Contact Contact NetApp technical support. for assistance.

#### Syslog Message

Failed to decompress a block on volume "%s" (vvbn=%llu pvbn=%llu). Details: %s.

#### Parameters

**vol** (STRING): Volume name.

**vvbn** (LONGINT): Virtual volume block number for which decompression failed.

**pvbn** (LONGINT): Physical volume block number for which decompression failed.

**details** (STRING): Additional information.

## wافل.cdemap events

### wافل.cdemap.repaired.metafile

#### Severity

ERROR

#### Description

This message occurs when an inconsistency is discovered in one of the metafile blocks. The inconsistency is automatically repaired upon detection.

#### Corrective Action

(None).

#### Syslog Message

Inconsistent block (%s) found in %s metafile on aggregate (%s). Repaired.

#### Parameters

**state** (STRING): A description of the metafile block that was found to be inconsistent.

**metafile\_name** (STRING): Name of the metafile.

**aggr** (STRING): Name of the aggregate.

## wافل.check events

## **wافل.صفء.صفء**

### **Severity**

ERROR

### **Description**

Generic wافل صفء message

### **Corrective Action**

(None).

### **Syslog Message**

%s

### **Parameters**

**msg** (STRING): The message from wافل صفء

## **wافل.صفء.لصفء.صفء**

### **Severity**

NOTICE

### **Description**

Wافل صفء failed to cleanup the lostnfound directory present in the metadir of the given volume.

### **Corrective Action**

(None).

### **Syslog Message**

Failed to cleanup lostnfound from the metadir of volume %s (reason: %s)

### **Parameters**

**vol** (STRING): The volume name

**err** (STRING): The error string

## **wافل.cleanup events**

### **wافل.cleanup.after.giveback**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs after Wافل® completes the post-giveback cleanup. The cleanup includes freeing memory resources allocated for the partner file system.

### **Corrective Action**

(None).



## Syslog Message

(None).

## Parameters

**delay** (LONGINT): Amount of time (in milliseconds) WAFL took to perform the post giveback cleanup. 1518 blocks

## wافل.cleanup.after.relocate

### Severity

INFORMATIONAL

### Description

This message occurs after WAFL® completes the aggregate relocation cleanup. The cleanup includes freeing memory resources allocated for the partner file system.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**aggrName** (STRING): Name of the aggregate.

**delay** (LONGINT): Amount of time (in milliseconds) WAFL took to perform the post aggregate relocation cleanup.

## wافل.cleanup.after.sendhome

### Severity

INFORMATIONAL

### Description

This message occurs after WAFL® completes the aggregate post-giveback cleanup. The cleanup includes freeing memory resources allocated for the partner file system.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**aggrName** (STRING): Name of the aggregate.

**delay** (LONGINT): Amount of time (in milliseconds) WAFL took to perform the aggregate post-giveback cleanup.

## wافل.clone events

## **waf1.clone.invalid.inherited.counters**

### **Severity**

ERROR

### **Description**

This message alerts about invalid inherited counters for the clone flexible volume.

### **Corrective Action**

(Call support).

### **Syslog Message**

clone inherited counters for clone flexible volume %s on aggregate %s are invalid. The clone volume will be offlined. It is recommended to online parent volume if not already onlined. Once the parent volume is onlined, the clone also be onlined. Contact technical support for assistance.

### **Parameters**

**clonename** (STRING): Clone flexible volume name.

**aggrname** (STRING): Aggregate name.

## **waf1.clone.no.basesnap**

### **Severity**

ERROR

### **Description**

This message alerts you about a missing clone flexible volume base Snapshot(tm) copy and strongly recommends that you take the clone flexible volume offline and destroy it.

### **Corrective Action**

Take the clone flexible volume offline immediately. If you know that the clone does not contain any valuable data, you can safely destroy and re-create it.

### **Syslog Message**

Base Snapshot copy for clone flexible volume %s on aggregate %s is missing. It is strongly recommended that you take the clone offline and destroy it. Failure to do so might lead to clone panic at any time.

### **Parameters**

**clonename** (STRING): Clone flexible volume name.

**aggrname** (STRING): Aggregate name.

## **waf1.clone.quota.error.ignored**

### **Severity**

ERROR

### **Description**

This message occurs when a clone op cannot fail due to quota error because either quota accounting is ignored by the user (e.g, sfsr) or the clone op cannot fail (eg, split), so it ignores the quota error it received during quota check.

## Corrective Action

None required.

## Syslog Message

Sis clone running in (vol='%s', volident='%s') ignored quota error '%u' during '%s' operation.

## Parameters

**vol** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**error\_code** (INT): A numerical error code indicating the reason why the quota check failed.

**op\_type** (STRING): Type of sis clone operation (create/split).

## waf.clone.snapSpcGuar

### Severity

INFORMATIONAL

### Description

This event is issued if an attempt is made to delete a snapshot in a clone volume that would require the space guarantee on the clone to exceed the available space in the aggregate. A clone volume initially shares one or more snapshots with its parent volume. This allows the space guarantee to be optimized, reducing the space requirement in the aggregate. If a shared snapshot is deleted then extra space is required in the aggregate to ensure the space guarantee for the volume can be enforced. If the freespace in the aggregate is too small to allow this then the snap delete is rejected.

### Corrective Action

This state can be corrected by performing one or more of the following actions:

- o Change the space guarantee on the volume so that less space is required in the aggregate.
- o Free space in the aggregate by shrinking or removing volumes.
- o Add space to the aggregate, e.g. by adding disks.

### Syslog Message

Snapshot '%s' in clone volume '%s%s%s' cannot be deleted, there is insufficient space in the aggregate to enforce the space guarantee (%s)

### Parameters

**snap\_name** (STRING): The snapshot name

**vol\_name** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**guarantee** (STRING): The current guarantee

## waf.clone.snapsplit.nospace

### Severity

NOTICE

### Description

This message occurs when there is not enough free space to complete a file clone split operation from the Snapshot(tm) copy. The clone file might become unreadable when the parent Snapshot copy is automatically deleted.

## Corrective Action

Increase the space available to the volume by either adding disks or by freeing space in the volume to allow the split operation to continue.

## Syslog Message

File clone split from Snapshot copy of '%s%s%s' cannot be completed due to lack of disk space.

## Parameters

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf1.clone.split.start.vol.fenced

### Severity

NOTICE

### Description

This event indicates that vol clone split start was called for a fenced volume after iron completion.

### Corrective Action

(None).

### Syslog Message

Volume clone split start called for fenced volume '%s%s%s'.

### Parameters

**volname** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.clone.unrecoverable

### Severity

ALERT

### Description

This message occurs during boot when a clone is marked unrecoverable due to its missing base Snapshot(tm) copy.

### Corrective Action

If you know that the clone does not contain any valuable data, you can safely destroy and re-create it. If it contains any valuable data, you can recover only the data which was written after the volume clone was created. Trying to access the data which was shared with the base snapshot (original data) will result in error.

### Syslog Message

Base Snapshot copy for clone %s on aggregate %s is missing. Marking it unrecoverable.

### Parameters

**clonename** (STRING): Clone name.

**aggrname** (STRING): Aggregate name.

## wafl.compaction events

### wafl.compaction.CP.value

#### Severity

INFORMATIONAL

#### Description

This message occurs when the aggregates compaction CP value is changed.

#### Corrective Action

(None).

#### Syslog Message

Compaction CP count in aggregate "%s" changed from %ld to %ld at aggregate's CP count %ld.

#### Parameters

**volname** (STRING): Volume name on which value was changed.

**oldval** (LONGINT): Old value of CP count.

**newval** (LONGINT): New value of CP count.

**fs\_cpcount** (LONGINT): CP count of aggregate fsinfo.

## wafl.compress events

### wafl.compress.cde.auto.enablement

#### Severity

NOTICE

#### Description

This message occurs when WAFL compression on container file level is automatically enabled on a volume.

#### Corrective Action

(None).

#### Syslog Message

WAFL auto adaptive compression automatically enabled in %s "%s".

#### Parameters

**volume\_identifier** (STRING): Volume identifier.

**volume** (STRING): Volume name.

### wafl.compress.cde.event

#### Severity

NOTICE

**Description**

This message occurs when WAFL compression on container file level is enabled or disabled on a volume.

**Corrective Action**

(None).

**Syslog Message**

WAFL container compression state changed in %s "%s" to "%s".

**Parameters**

**volume\_identifier** (STRING): Volume identifier.  
**volume** (STRING): Volume name.  
**message** (STRING): Whether enable or disable.

**waf1.compress.enable.failed****Severity**

INFORMATIONAL

**Description**

This event is issued when FAS compression can not be enabled on a volume. The following reasons might be causing the problem: FAS compression might not be supported on the software release you're running; the volume is not a flexible volume or it is a root volume or snaplock volume; the compression/de-compression scanner is already active; or a clone operation is already underway.

**Corrective Action**

(None).

**Syslog Message**

Compression could not be enabled on volume "%s" at this time. Check the volume configuration and retry.

**Parameters**

**vol** (STRING): Volume name.

**waf1.compress.fail.rootVol****Severity**

NOTICE

**Description**

This message occurs when an administrator attempts to enable compression on the root volume. Compression cannot be enabled on the root volume.

**Corrective Action**

(None).

**Syslog Message**

'%s' is the root volume and compression cannot be enabled on the root volume.

**Parameters**

**vol** (STRING): Volume name.

## waf1.compress.mark.cg.bad.hw

### Severity

ALERT

### Description

This message occurs when the FAS compression writes bad blocks for an inconsistent or corrupted compression group due to hardware errors.

### Corrective Action

Restore the corrupted file from the Snapshot copy, or contact NetApp technical support

### Syslog Message

FAS compression: Writing bad blocks, at fbn:%llu level:%u in %s inode (fileid:%d snapid:%u) in volume %s%s%s%s.

### Parameters

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**ino\_type** (STRING): Inode type

**fileid** (INT): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf1.compress.mark.cg.bad.sw

### Severity

ERROR

### Description

This message occurs when the FAS compression writes bad blocks for an inconsistent or corrupted compression group due to software errors.

### Corrective Action

Restore the corrupted file from the Snapshot copy, or contact NetApp technical support

### Syslog Message

FAS compression: Writing bad blocks, at fbn:%llu level:%u in %s inode (fileid:%d snapid:%u) in volume %s%s%s%s.

### Parameters

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**ino\_type** (STRING): Inode type

**fileid** (INT): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## **waf1.compress.scan.checkpoint.deleted**

### **Severity**

NOTICE

### **Description**

This event is generated when the FAS compression of existing data for volume deletes checkpoint.

### **Corrective Action**

(None).

### **Syslog Message**

vol compress: Checkpoint deleted for compression of existing data for volume "%s%s%s".

### **Parameters**

**vol** (STRING): Name of the compressed volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **waf1.convert events**

### **waf1.convert.directory.failed**

### **Severity**

ERROR

### **Description**

This message occurs when a directory cannot be safely converted to Unicode format. The directory contains files or subdirectories that must first be renamed.

### **Corrective Action**

Use the "volume file show-inode" command with the directory inode number and volume name information to find the path of the directory that could not be converted. There are some files or subdirectories within this parent directory that have names like "{123456789}". File names like these can make it unsafe for ONTAP® software to automatically convert a directory to Unicode format. To convert the directory, first rename the files or subdirectories in this parent directory that have names matching this pattern. To perform these renames, make sure that the volume has the "-convert-unicode" property set to "false", and then rename the content using an NFSv3 client.

### **Syslog Message**

Volume "%s:%s" was unable to convert directory %u to Unicode.

### **Parameters**

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volume** (STRING): Volume name that the directory is in.

**directory\_inum** (LONGINT): The Inode number (fileid) of the directory.



# wafl.cop events

## wafl.cop.flowcontrol.end

### Severity

NOTICE

### Description

This message occurs when the system is out of the flow control mode which was earlier initiated as a high number of volumes were receiving modify operations concurrently beyond a certain threshold.

### Corrective Action

(None).

### Syslog Message

The node is out of the flow control mode which was earlier initiated as a high number volumes were receiving modify operations concurrently beyond a certain threshold (%lu).

### Parameters

**high\_concurrency\_flowcontrol\_thres** (LONGINT): System defined threshold beyond which high-concurrency flow control mode is triggered.

## wafl.cop.flowcontrol.start

### Severity

NOTICE

### Description

This message occurs when the system is in a very high concurrency mode; i.e., when a high number of volumes receive modify operations concurrently beyond a certain threshold. The system is automatically getting into a flow control mode, where it may hold ops on certain volumes for longer time in order to reduce high-concurrency related impacts.

### Corrective Action

If the flow control situation is long-lived, and the delays are not acceptable, the concurrency level among the volumes need to be brought down by moving the volumes (preferably the ones encountering delays) out of this node.

### Syslog Message

Very high number of volumes (%lu) are receiving modify operations concurrently. The node is in a flow control mode (threshold: %lu). As a result few volumes may experience higher delay in latency.

### Parameters

**concurrent\_vvols\_cnt** (LONGINT): Number of volumes actively receiving modify operations.

**high\_concurrency\_flowcontrol\_thres** (LONGINT): System defined threshold beyond which high-concurrency flow control mode is triggered.

## wafl.cop.high.concurrency

## Severity

NOTICE

## Description

This message occurs when the system is in high-concurrency mode; i.e., when a high number of FlexVol® volumes receive modify operations concurrently during a monitoring interval. High concurrency might cause degradation in performance.

## Corrective Action

(None).

## Syslog Message

A high number of FlexVol volumes are receiving modify operations concurrently. Interval (ms): %lu, Aggregate: %s, Frequency: %lu out of %lu times, Threshold criteria: more than %u FlexVol volumes were concurrently receiving modify operations, Distribution: %s.

## Parameters

**monitoring\_interval** (LONGINT): Duration of monitoring interval in milliseconds.

**aggr** (STRING): Name of the aggregate.

**cp\_above\_vvol** (LONGINT): Number of Consistency Points (CPs) with FlexVol volumes participating above the vvol\_threshold value in this monitoring interval.

**total\_cp\_in\_interval** (LONGINT): Number of CPs that happened in the monitoring interval.

**vvol\_threshold** (INT): Minimum number of FlexVol volumes in a CP to call it a "highly concurrent CP".

**histogram** (STRING): A histogram [b1 : v1, b2 : v2, ..., b8 : v8], where "b" represents bin ID and "v" represents count of the corresponding bin ID.

# wافل.cp events

## wافل.cp.toolong

### Severity

ERROR

### Description

This message occurs when a WAFL® consistency point (CP) takes longer than 30 seconds. A CP lasting more than 30 seconds might cause client latency and potentially a client outage.

### Corrective Action

For additional information about this specific CP, see the kernel log file and the corresponding CM archive. For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "wافل.cp.toolong" keyword.

### Syslog Message

Aggregate %s experienced a long CP.

### Parameters

**aggrname** (STRING): Name of the aggregate.

**repeats** (INT): Number of consecutive long CPs on this volume.

**cp\_count** (INT): CP count on the volume corresponding to the long CP.

**total\_ms** (INT): Total time taken by the CP, in milliseconds.

**total\_dbufs** (LONGINT): Total number of dirty buffers in the CP.

**tetrises** (INT): Number of tetrises that were written across all raidgroups in volume.  
**blocks** (LONGINT): Number of blocks that were written across all raid groups in the volume.  
**rgid** (INT): ID of the raid group with the highest average tetris write time.  
**clean** (INT): Milliseconds to write allocate (clean) all normal files.  
**flush** (INT): Milliseconds to write all remaining blocks to disk.

## wafl.csm events

### wafl.csm.sessionStopped

#### Severity

INFORMATIONAL

#### Description

This message occurs when the Cluster Session Manager (CSM) session is stopped as part of the shutdown process.

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

(None).

## wafl.data events

### wafl.data.compaction.event

#### Severity

NOTICE

#### Description

This message occurs when data compaction functionality is enabled or disabled on a volume.

#### Corrective Action

(None).

#### Syslog Message

WAFL volume data compaction state changed in %s "%s" to "%s".

#### Parameters

**volume\_identifier** (STRING): Volume identifier.  
**volume** (STRING): Volume name.  
**message** (STRING): Whether enable or disable.

## wafl.delete events

## waf1.delete.bad.indirect

### Severity

NOTICE

### Description

This message occurs when the system encounters a lost indirect block while deleting a file. As a result, some space occupied by the file being deleted might not be reclaimed.

### Corrective Action

Run waf1iron to fix any space accounting discrepancies.

### Syslog Message

Lost indirect block encountered while deleting inode %d (level %d, vbn %llu) on %s %s%s%s%s. Contact NetApp technical support.

### Parameters

**ino** (INT): Inode number of the file.

**level** (INT): Indirection level of the block.

**vbn** (LONGINT): Block number of the block within the volume.

**type** (STRING): Type of the object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Name of the volume containing the file.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf1.delete.fenced.file

### Severity

NOTICE

### Description

This message occurs when the user attempts to delete a fenced file. Depending on the type of fence, reads and/or writes will not be permitted for the file.

### Corrective Action

Use the "volume file show-inode" command to confirm the file name that matches the inode, for example, "volume file show-inode -vserver 'vserver that fenced file is stored on' -volume 'volume that fenced file is stored on' -inode-number 'inode number of fenced file'". Clear the fenced status by using the "iofence clear 'file path name'" diagnostic privilege nodeshell command, and then the file can be deleted from the NFS/CIFS client.

### Syslog Message

The file with inode %d on volume %s cannot be deleted because it is fenced by %s operation.

### Parameters

**fid** (INT): Inode of the fenced file.

**vol** (STRING): Name of the volume that contains the fenced file.

**owner** (STRING): Name of the operation that fenced the file.

## waf.delete.specdev

### Severity

INFORMATIONAL

### Description

This message occurs when the system encounters a block or character device during zombie delete processing.

### Corrective Action

(None).

### Syslog Message

Block or character device encountered in zombie path: inode: %d, type: %d, size: %lld, block count: %llu. Possible inconsistency detected on %s %s%s%s%s.

### Parameters

**ino** (INT): Inode number.

**inotype** (INT): Inode type.

**size** (LONGINT): Size.

**blocks** (LONGINT): Block count.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**vendor** (STRING): Vendor contact information.

## waf.descend events

### waf.descend.fail

#### Severity

ERROR

#### Description

This message occurs when there is an error during descendency checking. This message can occur when accessing a directory in a Snapshot(tm) copy that is being deleted.

#### Corrective Action

This condition might be transient; try the operation again. If the problem persists, contact NetApp technical support.

#### Syslog Message

Descendency check failure: {x%x %d %d %d %d} parent of {x%x %d %d %d %d}. Contact NetApp technical support.

#### Parameters

**pfh0** (INT): Parent file handle's flags.

**pfh1** (INT): Parent file handle's snapid.

**pfh2** (INT): Parent file handle's fileid.

**pfh3** (INT): Parent file handle's generation.  
**pfh4** (INT): Parent file handle's FSID.  
**cfh0** (INT): Child file handle's flags.  
**cfh1** (INT): Child file handle's snapid.  
**cfh2** (INT): Child file handle's fileid.  
**cfh3** (INT): Child file handle's generation.  
**cfh4** (INT): Child file handle's FSID.  
**vendor** (STRING): Vendor contact information.

## wافل.dir.events

### wافل.dir.convert

#### Severity

NOTICE

#### Description

This message indicates that CIFS directory conversion is in progress and WAFL will not hang due to timeout. Directory conversion happens the first time a volume with 8.3-format names is accessed by a CIFS client.

#### Corrective Action

(None).

#### Syslog Message

WAFL %s directory conversion in progress.

#### Parameters

**type** (STRING): Type of directory operation (revert/convert).

### wافل.dir.convert.denied.2big

#### Severity

ERROR

#### Description

This message occurs when access to a directory is rejected because an attempted operation requires the directory to be converted to Unicode format and the directory is too large to be converted quickly. Directory conversion is required for a directory with 8.3-format names to be accessed by a CIFS or NFSv4 client.

#### Corrective Action

Use the "volume modify" command to enable the "convert\_ucose" option on the volume experiencing the rejected access, and then retry the access attempt. Note that this might result in system performance degradation or temporary loss of service while the conversion occurs. Search on "Converting existing directories to Unicode format" for more information.

#### Syslog Message

Access to large directory on volume MSID:%u, file ID:%u rejected for requiring conversion.

#### Parameters

**msid** (LONGINHEX): Volume master data set identifier (MSID) of the volume experiencing the rejected

access.

**fileid** (INT): File identifier of the directory experiencing the rejected access.

## **wافل.dir.depth.max**

### **Severity**

ERROR

### **Description**

This message occurs when a directory depth reaches its maximum depth limit. The default maximum depth is 2000.

### **Corrective Action**

Reduce the depth of the directory, by moving directory hierarchy to higher levels.

### **Syslog Message**

Directory depth for inum %d in volume %s%s exceeds the maximum directory depth limit. To fix this issue, please reduce the directory depth.

### **Parameters**

**fileid** (INT): Fileid of the directory for which directory depth reached its maximum limit.

**vol** (STRING): Volume in which directory depth reached its maximum limit.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier in cases where volume name itself is insufficient.

## **wافل.dir.duplicate.hash**

### **Severity**

NOTICE

### **Description**

This message occurs when a directory entry has a duplicate index hash for longname and dosname/altname is deleted from the directory.

### **Corrective Action**

(None).

### **Syslog Message**

An entry with duplicate index hash is found in directory inum %d in volume %s%s%s. Directory index clears to fix this issue.

### **Parameters**

**fileid** (INT): File ID of the directory.

**vol** (STRING): Volume name.

**app** (STRING): Application universal unique identifier(UUID).

**volident** (STRING): Unique volume identifier in cases where volume name itself is insufficient.

## **wافل.dir.fileld.err**

### **Severity**

ERROR

## Description

This event is issued when directory prefetch results in inconsistent file identifier data. The fh[0] to fh[4] parameters represent the encoded filehandle value.

## Corrective Action

If this error condition occurs, contact NetApp technical support for further assistance.

## Syslog Message

wafp\_prefetch\_dir\_inos: inconsistent fileid %d in %s directory {x%x %d %d %d %d}.

## Parameters

**fileid** (INT): The file identifier  
**ino\_type** (STRING): The inode type  
**fh0** (INT): fh[0]  
**fh1** (INT): fh[1]  
**fh2** (INT): fh[2]  
**fh3** (INT): fh[3]  
**fh4** (INT): fh[4]

## wafp.dir.link.max

### Severity

ERROR

### Description

This message occurs as a warning when a parent directory reaches its maximum link limit within a specified time interval and the no of parent directory reached maximum hard link limit reached the upper limit. A directory has only one parent directory, which counts as a hard link. All other hard links are from the directories contained within this directory. The link to the parent directory cannot be removed.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Reduce the hard link count of this directory by moving some of its subdirectories to another directory, or remove some of its subdirectories.

### Syslog Message

This is the last warning for failure to create a new sub-directory since parent directory with fileid %d in volume %s%s has reached its sub-directory limit.

### Parameters

**fileid** (INT): File ID of the parent directory that has reached its maximum hard link limit.  
**vol** (STRING): Name of the volume in which the parent directory resides.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory resides, when the volume name alone is insufficient.

## wafp.dir.link.reachingLimit

### Severity

ERROR



## Description

This message occurs when a parent directory surpasses 90% of its maximum hard link limit. The total hard link count for a directory is the number of its subdirectories plus other hard links to the directory.

## Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Move some of the parent directory's subdirectories to a different directory. Take care not to create too many new subdirectories in the parent directory that is reaching its limit until you reduce the hard link count.

## Syslog Message

The directory with fileid %d in volume %s%s%s is approaching the maximum link limit of %d. Reduce the number of links to the existing parent directory.

## Parameters

**fileid** (INT): File ID of the parent directory that has surpassed 90% of its maximum hard link limit.

**vol** (STRING): Name of the volume in which the parent directory resides.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory resides, when the volume name alone is insufficient.

**maxLink** (INT): Maximum number of hard links for the parent directory.

## wافل.dir.link.trap

### Severity

ALERT

### Description

This message occurs as a warning when a directory has reached its maximum hard link limit and the specified time interval is exceeded.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Reduce the hard link count of this directory by moving some of its subdirectories to another directory, or remove some of its subdirectories.

### Syslog Message

(None).

### Parameters

**fileid** (INT): File ID of the parent directory that has reached its maximum hard link limit.

**vol** (STRING): Name of the volume in which the parent directory resides.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory resides, when the volume name alone is insufficient.

**maxLink** (INT): Maximum number of hard links allowed.

## wافل.dir.link.warning

### Severity

ERROR

## Description

This event occurs as a warning when a parent directory reached its maximum hard link limit within a specified time interval and the no of parent directory reached maximum hard link limit is below the upper limit.

## Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. create new subdirectories in a different parent directory. Alternatively, move some of the current parent directory's subdirectories to a different directory.

## Syslog Message

A request to make a subdirectory in parent directory with fileid %d in volume %s%s could not be handled because the parent directory has reached its link limit; use a different parent directory or reduce the number of hard links by moving some of its subdirectories to a different directory.

## Parameters

**fileid** (INT): File ID of the parent directory that has reached its maximum hard link limit.

**vol** (STRING): Name of the volume in which the parent directory resides.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory resides, when the volume name alone is insufficient.

## wافل.dir.nt.max

### Severity

ERROR

### Description

This message occurs when an inode has reached its maximum number of NT streams.

### Corrective Action

Reduce the number of streams on the file.

### Syslog Message

Inode "%d" in volume "%s%s%s" has reached its maximum number of NT streams.

### Parameters

**fileid** (INT): File identifier.

**vol** (STRING): Name of the volume in which the inode resides.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the inode is present.

## wافل.dir.nt.warning

### Severity

NOTICE

### Description

This message occurs when an inode is nearing its maximum number of NT streams.

## Corrective Action

(None).

## Syslog Message

Inode "%d" in volume "%s%s%s" is nearing its maximum number of NT streams.

## Parameters

**fileid** (INT): File identifier.

**vol** (STRING): Name of the volume in which the inode resides.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the inode is present.

## wافل.dir.size.max

### Severity

EMERGENCY

### Description

This message occurs after a directory has reached its maximum directory size (maxdirsize) limit.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Reduce the number of files in the directory. If not possible, use the (privilege:advanced) option "volume modify -volume vol\_name -maxdir-size new\_value" to increase the maximum number of files per directory. However, doing so could impact system performance. If you need to increase the maximum directory size, contact NetApp technical support.

## Syslog Message

Directory size for fileid "%d" in volume "%s%s" reached the maxdirsize limit.

## Parameters

**fileid** (INT): File ID of the directory that has reached maximum size for directories in its volume.

**vol** (STRING): Volume in which the directory is present.

**app** (STRING): Application universal unique identifier(UUID).

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory resides, when the volume name alone is insufficient.

## wافل.dir.size.max.warning

### Severity

ERROR

### Description

This message occurs when the size of a directory has reached or surpassed a configured percentage (default: 90%) of its current maximum directory size (maxdirsize) limit, and the current maxdirsize is less than the default maxdirsize, which is 1% of total system memory.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Reduce the number of files in the directory. If not possible, use the (privilege:advanced) option "volume modify -volume vol\_name -maxdir-size new\_value" to increase the maximum number of files per directory. However, doing so could impact system performance. If you need to increase the maximum

directory size, contact technical support.

### Syslog Message

Directory size for file ID "%d" in volume "%s%s" is approaching the maximum directory size (maxdirsize) limit.

### Parameters

**fileid** (INT): File ID of the directory that has reached or is approaching the current maximum size for directories in its volume.

**vol** (STRING): Name of the volume in which the directory is located.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory is located, when the volume name alone is insufficient.

## wafldir.size.warning

### Severity

ALERT

### Description

This message occurs when the size of a directory surpasses a configured percentage (default: 90%) of its current maximum directory size (maxdirsize) limit.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Reduce the number of files in the directory. If not possible, use the (privilege:advanced) option "volume modify -volume vol\_name -maxdir-size new\_value" to increase the maximum number of files per directory. However, doing so could impact system performance. If you need to increase the maximum directory size, contact NetApp technical support.

### Syslog Message

Directory size for file ID "%d" in volume "%s%s" is approaching the maximum directory size (maxdirsize) limit.

### Parameters

**fileid** (INT): File ID of the directory that has surpassed 90% of current maximum size for directories in its volume.

**vol** (STRING): Name of the volume in which the directory is located.

**app** (STRING): Application universal unique identifier(UUID).

**volident** (STRING): Unique volume identifier (UUID) of the volume in which the parent directory is located, when the volume name alone is insufficient.

## wafldir.surrpair.filename

### Severity

ERROR

### Description

This message occurs as a warning when a file name with surrogate pair characters in UTF-16 encoding fails to be created in a parent directory.

## Corrective Action

To allow names with surrogate pairs to be created, use the following command: "setflag waf1\_reject\_surrogate\_pair 0". If the option needs to be set across reboots, set the bootarg 'waf1-accept-surrogate-pair?' to "true" at the LOADER prompt.

## Syslog Message

Failed to create a file with surrogate pair characters in the name in the directory %s.

## Parameters

**dir** (STRING): Parent directory.

# waf1.dirent events

## waf1.dirent.illegal.encoding

### Severity

ERROR

### Description

This message occurs when a user attempts to use a directory entry name that is not supported by the virtual server's language setting.

### Corrective Action

Use directory entry name with encoding that matches the virtual server's language setting.

### Syslog Message

The encoding of directory entry name %s doesn't match the virtual server's language code %s.

### Parameters

**filename** (STRING): The filename

**langcode** (STRING): Language code

# waf1.disable events

## waf1.disable.vol.guarantee

### Severity

NOTICE

### Description

This message occurs when the volume guarantee of a flexible volume is disabled due to aggregate space being physically overallocated.

### Corrective Action

Remount the aggregate after creating space by increasing the aggregate size or by deleting/moving some volumes to the other aggregates. Use the "storage aggregate show-space" command to view the current space usage of the aggregate. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot copies are deleted automatically to make some space when the aggregate is full.

## Syslog Message

Volume guarantee on volume "%s" is disabled after the aggregate "%s" becomes physically overallocated.

## Parameters

**vol** (STRING): Volume name whose guarantee is getting disabled.

**aggr** (STRING): The aggregate name on which the volume whose guarantee is getting disabled resides.

# wafldosname events

## wafldosname.repair.fail

### Severity

ERROR

### Description

This message occurs when the system cannot correct duplicate dosnames present in a directory.

### Corrective Action

From the NFS client, check the directory for file or subdirectory names like {}. Rename these, and then run (privilege diag) the "wafldosname" node shell command on the directory.

## Syslog Message

Dosname repair failed for file ID: %d volume %s:%s.

## Parameters

**fileid** (INT): Directory file ID where dosname repair failed.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volume** (STRING): Volume name.

## wafldosname.repair.success

### Severity

NOTICE

### Description

This message occurs when the "wafldosname" command is completed successfully.

### Corrective Action

(None).

## Syslog Message

"wafldosname" command successfully completed on volume %s:%s. Total files corrected: %d.

## Parameters

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's Vserver.

**volume** (STRING): Volume name.

**total** (INT): Total number of directory entries changed.

## wafl.enc events

### wafl.enc.aggrMFRecreated

#### Severity

NOTICE

#### Description

This message occurs when the system detects a corrupted WAFL aggregate encryption keys metafile and creates a corrected replacement.

#### Corrective Action

(None).

#### Syslog Message

"Recreated aggregate %s encryption keys metafile."

#### Parameters

**aggr** (STRING): Aggregate name.

### wafl.enc.aggrMFRecreateFail

#### Severity

ALERT

#### Description

This message occurs when the system detects a corrupted WAFL aggregate encryption keys metafile and is unable to create a corrected replacement. The system might not be able to mount volumes.

#### Corrective Action

Contact NetApp technical support.

#### Syslog Message

"Failed to recreate aggregate %s encryption key metafile."

#### Parameters

**aggr** (STRING): Aggregate name.

## wafl.encrypted events

### wafl.encrypted.aggr.created

#### Severity

NOTICE

#### Description

This message occurs when a NetApp® Aggregate Encryption (NAE) aggregate is created.

## Corrective Action

(None).

## Syslog Message

NAE aggregate "%s" is created.

## Parameters

**Aggregate** (STRING): Name of the aggregate.

**KeyID\_1** (STRING): Aggregate key ID of key1.

**KeyID\_2** (STRING): Aggregate key ID of key2.

**KeyID\_3** (STRING): Aggregate key ID of key3.

**KeyID\_4** (STRING): Aggregate key ID of key4.

## waf1.encrypted.aggr.modified

### Severity

NOTICE

### Description

This message occurs when an aggregate is converted to a NetApp® Aggregate Encryption (NAE) aggregate.

### Corrective Action

(None).

### Syslog Message

Aggregate "%s" is converted to a NAE aggregate.

### Parameters

**Aggregate** (STRING): Name of the aggregate.

**KeyID\_1** (STRING): Aggregate key ID of key1.

**KeyID\_2** (STRING): Aggregate key ID of key2.

**KeyID\_3** (STRING): Aggregate key ID of key3.

**KeyID\_4** (STRING): Aggregate key ID of key4.

## waf1.encryption events

### waf1.encryption.unsupported

#### Severity

ALERT

#### Description

This message occurs when the system takes offline an encrypted virtual volume because encryption became unsupported.

#### Corrective Action

Reenable support for encryption by setting the bootarg

"bootarg.disable.volume.softwareencryption.support" to "false", reboot the system and then bring online the encrypted volume.



## Syslog Message

Keeping encrypted volume "%s" (UUID: %s) offline because encryption has become unsupported.

## Parameters

**uuid** (STRING): UUID of the volume.

**vol** (STRING): Volume name.

# wafL.err events

## wafL.err.L1.compacted.block

### Severity

NOTICE

### Description

This message occurs when a level 1 indirect block with given block number is not found in a compacted block.

### Corrective Action

(None).

## Syslog Message

Block number 0x%llx is not found in the compacted block.

## Parameters

**vvbn** (LONGINT): Volume block number.

# wafL.extent events

## wafL.extent.size.reset

### Severity

NOTICE

### Description

This message occurs when the volume extent option is set and WAFL® encounters an invalid extent size. This message indicates that the extent size was reset to a valid default value.

### Corrective Action

(None).

## Syslog Message

Extent size for volume %s was reset to %d.

## Parameters

**vol** (STRING): Name of the volume.

**extent\_size** (INT): New extent size.

# wapl.failed events

## wapl.failed.mount

### Severity

EMERGENCY

### Description

This message occurs when the ONTAP® kernel cannot mount a volume because of file system inconsistency.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

WAPL inconsistent: Mount failed for volume %s%s%s%s.

### Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wapl.failed.mount.bad.fsid

### Severity

EMERGENCY

### Description

This message occurs when DATA ONTAP® cannot mount a volume because of a hardware failure in the disk subsystem. The failure invalidates and clears the unique identifier (FSID) that was assigned to the volume at the time of creation, resulting in an invalid FSID value of 0 and making the volume inaccessible for clients. The volume will be left in the offline state without any data loss.

### Corrective Action

Check for disk failures and replace any failed disks. A subsequent mount call for the volume, either from the boot path or an admin level online request, would be able to read a valid FSID now, thus solving this problem. For more information or assistance, contact NetApp technical support.

### Syslog Message

Invalid file system identifier (FSID): Mount failed for volume %s%s%s%s.

### Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name itself is insufficient.

# waf.fc events

## waf.fc.recl.csm.disconnect

### Severity

EMERGENCY

### Description

This message occurs when a FlexCache® origin node detects that its FlexCache volume has disconnected while a lock reclaim operation is in progress.

### Corrective Action

Verify that the intercluster LIFs on both nodes are operational.

### Syslog Message

FlexCache origin node with UUID %s detected potential disconnect with FlexCache with UUID %s while processing lock reclaims for the volume with MSID %u hosted on Vserver with UUID %s.

### Parameters

**localbladeuuid** (STRING): Universally unique identifier (UUID) of the origin node.

**bladeuuid** (STRING): UUID of the FlexCache Cache node.

**msid** (LONGINT): Volume master data set identifier (MSID) of the FlexCache volume.

**vserveruuid** (STRING): UUID of the FlexCache Vserver.

# waf.fgmember events

## waf.fgmember.rollback

### Severity

ERROR

### Description

This message occurs when a member of a FlexGroup experiences an unexpected rollback in its state due to a loss of the NVRAM log for the member. This indicates that the volume might have inconsistencies between members.

### Corrective Action

Run the fgiron tool against the FlexGroup to ensure consistency.

### Syslog Message

FlexGroup %u member %u/%u detected a rollback.

### Parameters

**FlexGroupMSID** (LONGINTEX): Volume master data set identifier (MSID) of the FlexGroup that experienced a rollback due to NVRAM loss.

**FlexGroupMemberMSID** (LONGINTEX): Volume master data set identifier (MSID) of the FlexGroup member that experienced a rollback due to NVRAM loss.

**FlexGroupMemberDSID** (LONGINTEX): Volume data set identifier (DSID) of the FlexGroup member that experienced a rollback due to NVRAM loss.

# wafl.filecgc events

## wafl.fileCG.vvol.size.exceeded

### Severity

ERROR

### Description

This message occurs when a volume with file-based compression(either secondary compression or 8k/16K/32K adaptive compression) have nominal size greater than 300TB.

### Corrective Action

(None).

### Syslog Message

Nominal size of the "%s" volume exceeded the 300TB limit.

### Parameters

**vvolname** (STRING): Name of the volume.

# wafl.fs events

## wafl.fs.corrupt.fr.pct.changed

### Severity

NOTICE

### Description

This message occurs during volume mount, when the fractional reservation value for a volume is found to be corrupted. This is necessary because the storage system software only supports fractional reservation values of 0% and 100%. The guarantee type setting of a volume determines what the correct fractional reservation must be. For volumes with a guarantee type of VOLUME or NONE the value must be corrected to be 0%. For volumes with a guarantee type of FILE the value must be set to 100%.

### Corrective Action

(None).

### Syslog Message

The corrupted fractional reservation value for volume '%s%s%s' with guarantee type '%s' has been changed from %d%% to %d%%.

### Parameters

**vol** (STRING): Name of the volume whose fractional reservation value was changed.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**guarantee\_type** (STRING): Guarantee type of the volume.

**old\_fr\_pct** (INT): Old fractional reservation value of the volume.

**new\_fr\_pct** (INT): New fractional reservation value of the volume.

## wafl.fs.file.guarnt.upgraded

### Severity

ERROR

### Description

This message occurs during a volume upgrade, when the space guarantee style for a volume is changed to either "none" or "volume" (guarantee is disabled). This is necessary because Data ONTAP® no longer supports "file" space guarantee style for flexible volumes.

### Corrective Action

Use the "volume show -volume <volume-name> -fields space-guarantee" command to display the volume's new space guarantee style. To modify the style to "volume", ensure that there is enough space in the aggregate, and then enter the "volume modify -vserver <vserver-name> -space-guarantee volume <volume-name>" command. To modify the style to "none" enter the "volume modify -vserver <vserver-name> -space-guarantee none <volume-name>" command.

### Syslog Message

The space guarantee type for volume '%s%s%s' has been upgraded from file to %s.

### Parameters

**vol** (STRING): Name of the volume whose file guarantee was changed.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**new\_guarantee\_type** (STRING): New space guarantee type chosen for the volume.

## wafl.fs.upgrade.fr.pct.changed

### Severity

NOTICE

### Description

This message occurs during a volume upgrade, when the fractional reservation value for a volume is checked and corrected. This is necessary because the storage system software only supports fractional reservation values of 0% and 100%. The guarantee type setting of a volume determines what the fractional reservation value must be. For volumes with a guarantee type of VOLUME or NONE the value must be set to 0%. For volumes with a guarantee type of FILE the value must be set to 100%.

### Corrective Action

(None).

### Syslog Message

The fractional reservation value for volume '%s%s%s' with guarantee type '%s' has been changed from %d%% to %d%%.

### Parameters

**vol** (STRING): Name of the volume whose fractional reservation value was changed.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**guarantee\_type** (STRING): Guarantee type of the volume.

**old\_fr\_pct** (INT): Old fractional reservation value of the volume.

**new\_fr\_pct** (INT): New fractional reservation value of the volume.

# wafil.gd events

## wafil.gd.convert.scan.running

### Severity

NOTICE

### Description

This message occurs when converting a file to use granular data fails because a scanner is currently running on the volume. The conversion cannot be completed until the scan is complete or filesystem inconsistencies can occur. Granular data conversion is required for non-disruptive file moves that occur as part of volume rebalancing.

### Corrective Action

Retry the operation after some time to allow the scan to complete. If the problem persists, contact NetApp technical support for assistance.

### Syslog Message

Granular data conversion failed for fileid "%d" on volume "%s%s%s%s" because scan "%s" is currently running on the volume.

### Parameters

**fileid** (INT): File identifier.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier that uniquely identifies the volume in cases where the volume name itself is insufficient.

**scantype** (STRING): Type of scan.

## wafil.gd.convert.unknown.pvbn

### Severity

NOTICE

### Description

This message occurs when converting a file to use granular data fails because the file contains unresolved physical block numbers (PVBNs). Unresolved PVBNs occur when you move a volume to a different aggregate. PVBNs are resolved by a background process following the volume move.

### Corrective Action

Retry the operation after some time to allow the background process to complete.

### Syslog Message

Granular data conversion failed for fileid "%d" on volume "%s%s%s%s" because the inode contains unresolved PVBNs.

### Parameters

**fileid** (INT): File identifier.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier that uniquely identifies the volume in cases where the volume name itself is insufficient.

## wafl.giveback events

### wafl.giveback.create.fail

#### Severity

ERROR

#### Description

This message occurs when the system cannot perform a giveback due to vol create, and must abort giveback.

#### Corrective Action

A volume is being created. Wait for volume creation to finish and then retry giveback.

#### Syslog Message

%s %s%s%s%s is being created; canceling giveback.

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

### wafl.giveback.fence.fail

#### Severity

ERROR

#### Description

This message occurs when the system cannot perform a giveback operation because a fence is set on one of the volumes in the aggregate. The giveback operation is aborted.

#### Corrective Action

A volume in the aggregate is currently fenced for group operations. Retry giveback after any group SnapRestore® or mirror transfers are complete. If the problem persists, contact technical support.

#### Syslog Message

%s %s%s%s%s has a fence set; canceling giveback.

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.giveback.iron.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform a giveback due to wafiron, and must abort giveback.

### Corrective Action

(Call support).

### Syslog Message

%s %s%s%s%s has vol wafiron start underway; canceling giveback.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafl.giveback.sp.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform a giveback operation because a secure purge operation is being performed on one of the volumes in the aggregate. The giveback operation is aborted.

### Corrective Action

A volume in the aggregate is undergoing a secure purge operation. Retry giveback after this operation is complete, or after you abort the operation by using the "volume encryption secure-purge abort" command. If the problem persists, contact NetApp technical support.

### Syslog Message

A secure purge operation is in progress on %s %s%s%s. Canceling giveback.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.giveback.statusvol.fail

### Severity

ERROR



## Description

This message occurs when the system cannot perform a giveback due to the volume being used to host wafiron status files.

## Corrective Action

Use "aggr wafiron show" command to monitor the status of wafiron and retry giveback after the completion of wafiron that uses the status files hosting volume.

## Syslog Message

%s %s%s%s%s is being used by wafiron to host status files. Canceling giveback.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafl.giveback.vetoedByNvsave

### Severity

ERROR

### Description

This message occurs when a high-availability (HA) giveback is vetoed because there are NVSave files for an aggregate or a flexible volume that must be replayed before the aggregate or flexible volume can be given back.

### Corrective Action

Retry the giveback operation after the NVSave files either are replayed (by bringing the aggregate or flexible volume online) or are deleted from the root volume. Successful replay, for example, is indicated by the EMS message 'wafl.nvsave.replaying.pass'.

### Syslog Message

Giveback vetoed because NVSave files are present for %s FSID %d name %s%s%s. UUID is %s.

### Parameters

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

**fsid** (INT): File system identifier.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

**UUID** (STRING): UUID of the volume.

## wafl.gs events

### wafl.gs.metafile.load.failed

#### Severity

ERROR

## Description

This message occurs when waf1 is unsuccessful to check and load Granular Set metafile for the volume.

## Corrective Action

Issue volume offline followed by volume online and confirm this error message does not happen again. For more information or assistance, contact NetApp technical support.

## Syslog Message

Waf1 failed to load metafile that stores waf1 copy of SnapMirror CG.

## Parameters

**volume** (STRING): Volume name.

**error** (INT): Error code.

## waf1.gs.modify.op.failed

### Severity

NOTICE

### Description

This message occurs when WAFL® prohibits loading a file for a modify operation on a read-only consistency group.

### Corrective Action

NONE.

### Syslog Message

Failed to load file ID %u on volume %s%s%s%s FSID:%u for a modify operation because it is part of a read-only sync relationship.

### Parameters

**fileid** (INT): File identifier.

**volowner** (STRING): Owner of the volume.

**volname** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient. Consists of the Vserver name and UUID.

**fsid** (INT): File system ID of the volume.

## waf1.gs.remove.rename.op.failed

### Severity

NOTICE

### Description

This message occurs when WAFL® prohibits a rename operation or an unlink operation on a file part of a consistency group.

### Corrective Action

NONE.

## Syslog Message

Failed to load file ID %u on volume %s%s%s%s FSID:%u for a rename operation or an unlink operation because it is part of a sync relationship.

## Parameters

**fileid** (INT): File identifier.  
**volowner** (STRING): Owner of the volume.  
**volname** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient. Consists of the Vserver name and UUID.  
**fsid** (INT): File system ID of the volume.

## waf.gs.resync.gen.check.failed

### Severity

NOTICE

### Description

This message occurs when WAFL® prohibits modification to a file that is part of a consistency group and the operation does not have the right generation enforcement parameters. This can occur if the operation happened before the SnapMirror® relationship was established.

### Corrective Action

NONE.

## Syslog Message

Unexpected operation on consistency group file ID %u on volume %s%s%s%s FSID:%u. This can occur if the operation happened before the SnapMirror relationship was established.

## Parameters

**fileid** (INT): File identifier.  
**volowner** (STRING): Owner of the volume.  
**volname** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient. Consists of the Vserver name and UUID.  
**fsid** (INT): File system ID of the volume.

## wafl.hya events

### wafl.hya.aggr.policy.init

### Severity

INFORMATIONAL

### Description

This message occurs when the default caching policies are to be configured on all volumes contained in this aggregate.

**Corrective Action**

(None).

**Syslog Message**

Attempting to configure the default caching policies on all volumes in the aggregate '%s %s'.

**Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**waf1.hya.bc.add.cbi****Severity**

INFORMATIONAL

**Description**

This message occurs when a hybrid aggregate adds one counter file block (CBI - counter block index).

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**aggr** (STRING): Aggregate name.

**aggr\_uuid** (STRING): Universal Unique ID (UUID) of the aggregate.

**volume** (STRING): Volume name.

**volume\_uuid** (STRING): UUID of the volume.

**btid** (INT): Buftree ID.

**cbi** (INT): CBI value.

**volume\_info** (STRING): Volume type information.

**waf1.hya.bc.del.flexvol****Severity**

INFORMATIONAL

**Description**

This message occurs when a hybrid aggregate block accounting module receives a FlexVol® deletion notification.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**aggr** (STRING): Aggregate name.

**aggr\_uuid** (STRING): UUID of the aggregate.

**volume** (STRING): Volume name.  
**volume\_uuid** (STRING): UUID of the volume.  
**volume\_info** (STRING): Volume type information.  
**btid** (INT): Buftree ID.  
**cbi** (INT): CBI (counter block index) value.  
**ctr0** (LONGINT): Counter 0, total used blocks by this buftree ID at the time of the delete event.

## waf1.hya.bc.free.cbi

### Severity

INFORMATIONAL

### Description

This message occurs when a hybrid aggregate block accounting module frees a CBI (counter block index).

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**aggr** (STRING): Aggregate name.  
**aggr\_uuid** (STRING): UUID of the aggregate.  
**btid** (INT): Buftree ID.  
**cbi** (INT): CBI value.  
**ctr0** (LONGINT): Counter 0, total used blocks by this buftree ID at the time of the free event.

## waf1.hya.cache.ret.policy

### Severity

INFORMATIONAL

### Description

This message occurs when the cache retention priority policy for a volume in a hybrid aggregate has been modified.

### Corrective Action

(None).

### Syslog Message

Cache retention priority policy on volume '%s%s%s%s' has been modified to %s.

### Parameters

**owner** (STRING): Owner of the volume.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.  
**cache\_retention\_priority** (STRING): Cache retention priority policy configured on the volume.

## **waf1.hya.policy.init.skipped**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when cache policies could not be configured because of an invalid volume state.

### **Corrective Action**

(None).

### **Syslog Message**

Default caching policies were not set on volume '%s%s%s'.

### **Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## **waf1.hya.rcache.policy**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the read caching policy for a volume in a hybrid aggregate has been modified.

### **Corrective Action**

(None).

### **Syslog Message**

Read caching policy on volume '%s%s%s%s' has been modified to '%s'.

### **Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**read\_policy** (STRING): Read caching policy configured on the volume.

## **waf1.hya.sizer.warning**

### **Severity**

NOTICE

### **Description**

This message occurs when Flash Pool (tm) AWA (Automatic Workload Analyzer) encounters an abnormal condition such as low sampling rate or low memory resource.

## Corrective Action

Contact NetApp technical support for assistance in resolving this issue.

## Syslog Message

(None).

## Parameters

**aggr** (STRING): Aggregate name.

**aggr\_uuid** (STRING): UUID of the aggregate.

**vol** (STRING): Name of the volume.

**event\_info** (STRING): Flash Pool (tm) Sizer warning information.

## waf1.hya.wcache.policy

### Severity

INFORMATIONAL

### Description

This message occurs when write caching policy for a volume in a hybrid aggregate has been modified.

### Corrective Action

(None).

### Syslog Message

Write caching policy on volume '%s%s%s%s' has been modified to '%s'.

### Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**write\_policy** (STRING): Write caching policy configured on the volume.

## waf1.i2p events

### waf1.i2p.entries.mismatch

#### Severity

NOTICE

#### Description

This message occurs when the ONTAP® kernel detects that an i2p metafile on one of the snapshot has more entries than necessary for an inode.

#### Corrective Action

(None).

#### Syslog Message

There are more entries in the i2p metafile than there should be in volume %s%s%s%s, for inode %d in snapshot %d.

## Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases where the volume name itself is insufficient.

**fileid** (INT): Identifying number for the suspect inode.

**snapid** (INT): The snapshot identifier.

## wافل.i2p.fix.parent.pointer

### Severity

NOTICE

### Description

This message occurs when an inode-to-parent (i2p) scan detects and rectifies an incorrect parent inode number or cookie in an inode while fixing i2p data in a converted directory.

### Corrective Action

(None).

### Syslog Message

%s %s%s%s%s: Inode %u, fixing parent pointer, was: (parent: %u, cookie: 0x%x).

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**inode** (INT): Inode number.

**old\_parent\_inode** (INT): Existing parent inode number.

**old\_cookie** (INT): Existing parent cookie.

## wافل.ialloc events

### wافل.ialloc.is.disabled

#### Severity

NOTICE

#### Description

This message occurs when an inode allocation in a stripe is disabled.

#### Corrective Action

(None).

#### Syslog Message

Inode allocation in stripe is disabled for %s inode space.



## Parameters

**ino\_type** (STRING): Inode type.

# wafl.inaccessible events

## wafl.inaccessible.userdata

### Severity

ERROR

### Description

This message occurs when an access request is received for inaccessible user data (area blocked by wafiron).

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Recover the file from a backup. For a SnapMirror® target, rebuild from the source or resynchronize. If the block is in a Snapshot copy, delete the Snapshot copy. Contact NetApp technical support for more options.

### Syslog Message

Inaccessible user data block at FBN %llu Level %u in %s inode (file ID:%s Snapshot copy ID:%u file\_type:%d disk\_flags:0x%x) in volume %s%s%s%s.

## Parameters

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**file\_space** (STRING): File space.

**fileid** (STRING): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**type** (INT): File type.

**flags** (LONGINT): Inode disk flags.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

# wafl.incomplete events

## wafl.incomplete.i2p

### Severity

ERROR

### Description

This event occurs when the WAFL® kernel detects incomplete inode to path information for an inode on disk. This will cause the inode to path database to be reconstructed.

### Corrective Action

No immediate action is required. The Inode to Path database will be reconstructed in the background. Please contact NetApp technical support or disable i2p using the volume option no\_i2p.

## Syslog Message

Incomplete Inode to Path information detected in volume %s%s%s%s. I2P data will be reconstructed.

## Parameters

**owner** (STRING): The owner of the volume  
**vol** (STRING): The name of the volume  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.  
**func** (STRING): The internal function name  
**fileid** (INT): The identifying number for the suspect inode  
**snapid** (INT): Snapshot identifier (0) - active file system (non-zero) - refers to snapshot  
**reason** (STRING): An indication of the nature of the event.

# wafl.incons events

## wafl.incons.compr.metadata

### Severity

ERROR

### Description

This message occurs when the WAFL® kernel detects inconsistent compression metadata for a compression group.

### Corrective Action

If the Snapshot copy ID (snapid) is 0, the inconsistency is in the active file system and you need to run wafiron. To run wafiron, first find the aggregate the volume belongs to by running the "volume show -volume [ vol\_name ]" command. Then enter diag mode by using the "set diagnostic" command, and start wafiron by running the command: "storage aggregate wafiron start -aggregate [ aggr\_name ]". If the snapid is not 0, then run the "snapshot show -volume [ vol-name ] -instance" command, and then delete the Snapshot copy whose Snap ID field matches the given snapid.

## Syslog Message

WAFL inconsistent: bad compression indirect found (level:%u compression group first vbn:%llu) in %s inode (fileid:%d snapid:%u) in volume %s%s%s%s.

## Parameters

**level** (INT): Metadata level.  
**cg\_first\_vbn** (LONGINT): First block of the compression group in which the inconsistency was found.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): Identifying number for the suspect inode.  
**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy  
**owner** (STRING): Owner of the volume.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## wafl.incons.compressed.buf

### Severity

ERROR

## Description

This message occurs when the WAFL® kernel detects an inconsistent compressed block on a disk.

## Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Restore the inconsistent file from the Snapshot copy.

## Syslog Message

WAFL inconsistent: bad compressed block at VBN %llu (vvbn:%llu fbn:%llu level:%u compression group first vbn:%llu compression group size %d) in %s inode (fileid:%d snapid:%u) in volume %s%s%s%s.

## Parameters

**vbn** (LONGINT): Volume block number.

**vvbn** (LONGINT): Virtual block number.

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**cg\_first\_vbn** (LONGINT): First block of the compression group in which the inconsistency was found.

**cg\_size** (INT): Number of allocated blocks in the compression group.

**ino\_type** (STRING): Inode type.

**fileid** (INT): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## wafl.incons.objslothdr.buf.fixup.required

### Severity

INFORMATIONAL

### Description

This message occurs when number of inconsistent objslothdr entries are more than threshold on a FabricPool aggregate with unified format enabled.

### Corrective Action

For further assistance, contact NetApp technical support.

### Syslog Message

Number of inconsistent objslothdr entries(%d) are more than threshold. For further assistance, contact NetApp technical support.

### Parameters

**invalid\_entry\_count** (INT): Number of invalid entries present in slothdr metafile.

**aggregate\_name** (STRING): Aggregate name.

## wafl.incons.snap.metadata

### Severity

ERROR

## Description

This message occurs when a Snapshot(tm) copy or its metadata is found to be inconsistent.

## Corrective Action

Delete the inconsistent Snapshot copy by using the "volume snapshot delete" command.

## Syslog Message

WAFL inconsistent: Inconsistent Snapshot copy found in volume %s%s%s%s with Snapshot copy ID %d.  
Reason: %s.

## Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**snapid** (INT): Snapshot copy identifier.

**reason** (STRING): Reason the Snapshot copy is bad.

## wافل.incons.tlog

### Severity

ERROR

### Description

This message occurs when an object cannot be written to the object store due to a corrupted data block within one of the metafiles.

### Corrective Action

Contact NetApp technical support for further assistance.

### Syslog Message

Inconsistency found in the staging metafile in FabricPool aggregate %s. Skip writing object to object store due to bad block (pvbn %llu).

### Parameters

**aggrname** (STRING): Name of the aggregate.

**pvbn** (LONGINT): Unloadable block number that resulted in skipping a write of the corresponding object to the object store.

## wافل.incons.userdata.vol

### Severity

ALERT

### Description

This message occurs when the WAFL® kernel discovers an inconsistent volume with bad user data blocks.

### Corrective Action

The volume has an inconsistent user data block. Use the fileid and volume name information to run the 'inodepath' command. The 'inodepath' command will output the path name for the fileid. Recover the file from a backup. For a SnapMirror® target, rebuild from the source or resynchronize. If the block is in a Snapshot(tm) copy, delete the Snapshot copy. Then contact NetApp technical support.

## Syslog Message

WAFL inconsistent: %s %s%s%s%s has an inconsistent user data block. Note: Any new Snapshot copies might contain this inconsistency.

## Parameters

**type** (STRING): Container type (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**vendor** (STRING): Vendor name.

# waf.inconsistent events

## waf.inconsistent.acl

### Severity

ERROR

### Description

This message occurs when the ONTAP® kernel detects an inconsistent WAFL® Access Control List (ACL) on disk. Access to this ACL is disabled until you repair it by running wafiron.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

WAFL inconsistent: bad ACL (xfileid:%u snapid:%u) for inode (fileid:%u) in volume %s%s%s%s.

## Parameters

**xfileid** (INT): Identifying number for the suspect ACL.

**snapid** (INT): Snapshot(tm) copy identifier (0) - active file system (non-zero) - refers to Snapshot copy.

**fileid** (INT): Identifying number for the inode.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## waf.inconsistent.acl.data

### Severity

ERROR

### Description

This message occurs when Data ONTAP detects inconsistent Access Control List (ACL) data on disk.

### Corrective Action

The volume has inconsistent Access Control List (ACL) data. Use the "volume file show-inode" command with the file ID and volume name information to find the file path. .Rebuild ACL for the file.

## Syslog Message

WAFL inconsistent: bad ACL data (acl\_fileid:%u acl\_snapid:%u) for inode (fileid:%u snapid:%u file\_type:%d disk\_flags:0x%x) in volume %s%s%s%s%s.

## Parameters

**acl\_fileid** (INT): Identifying number for the ACL inode

**acl\_snapid** (INT): Snapshot(tm) copy identifier for ACL inode: (0) - active file system (non-zero) - refers to Snapshot copy

**fileid** (INT): Identifying number for the inode that owns the ACL.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**type** (INT): File type.

**flags** (LONGINT): Inode disk flags.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

**inotype** (STRING): The inode space (public or private)

## wafl.inconsistent.compacted.block

### Severity

NOTICE

### Description

This message occurs when an inconsistency is detected in a compacted block.

### Corrective Action

(None).

## Syslog Message

Inconsistent compacted block %llu in %s "%s". %s

## Parameters

**pvbn** (LONGINT): Volume block number.

**volume\_identifier** (STRING): Volume identifier.

**volume** (STRING): Volume name.

**message** (STRING): Type of inconsistency in compacted block.

## wafl.inconsistent.container

### Severity

ERROR

### Description

This message occurs when the WAFL® kernel detects a block in a FlexVol® volume that is inaccessible. Access to this block is disabled until you repair it by running wafliron.

### Corrective Action

The volume is inconsistent. Run wafliron as soon as possible to clear this inconsistency. For a SnapMirror® target, rebuild from the source or resynchronize. If the inode is in a Snapshot(tm) copy, delete the Snapshot copy, then contact NetApp technical support.

## Syslog Message

WAFL inconsistent: Volume %s%s%s%s has unreachable block at vvbn %llu (pvbn %llu).

## Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**vvbn** (LONGINT): Virtual block number.

**pvbn** (LONGINT): Physical block number.

## wافل.inconsistent.dir

### Severity

ALERT

### Description

This event occurs when the WAFL® kernel detects an inconsistent directory on disk. Access to this directory is disabled until you repair it by running wافلiron.

### Corrective Action

The volume is inconsistent. Run wافلiron as soon as possible to clear this inconsistency; for a snapmirror target, rebuild from source or resync or if inode is in a snapshot, delete the snapshot. Then contact NetApp technical support.

## Syslog Message

WAFL inconsistent: %s: bad directory ( %s fileid:%u, snapid: %u, type:%u, size:%u) in volume %s%s%s%s.

## Parameters

**func** (STRING): The internal function name

**ino\_type** (STRING): The inode type

**fileid** (INT): The identifying number for the suspect inode

**snapid** (INT): Snapshot identifier (0) - active file system (non-zero) - refers to snapshot

**type** (INT): The type of inode that's supposed to be a directory of some type

**size** (INT): The possibly corrupt size of the suspect inode.

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wافل.inconsistent.dirent

### Severity

EMERGENCY

### Description

This event occurs when an inconsistent directory entry has been detected. Access to this directory entry is disabled.

### Corrective Action

If this error condition occurs, contact NetApp technical support for further assistance.

## Syslog Message

WAFL inconsistent: %s: bad entry {x%x %d snapid: %d %d %d} <%. \*s> in %s directory {x%x %d snapid: %d %d %d} in volume %s%s%s%s.

## Parameters

**func** (STRING): The internal function name  
**fh0** (INT): fh[0]  
**fh1** (INT): fh[1] - snapid (0) - active file system (non-zero) - refers to snapshot  
**fh2** (INT): fh[2]  
**fh3** (INT): fh[3]  
**fh4** (INT): fh[4]  
**name** (STRINGARRAY): The name of the directory entry  
**ino\_type** (STRING): The inode type  
**dfh0** (INT): dfh[0]  
**dfh1** (INT): dfh[1] - snapid in filehandle (0) - active file system (non-zero) - refers to snapshot  
**dfh2** (INT): dfh[2]  
**dfh3** (INT): dfh[3]  
**dfh4** (INT): dfh[4]  
**owner** (STRING): The owner of the volume  
**vol** (STRING): The name of the volume  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wافل.inconsistent.ino

### Severity

ERROR

### Description

This message occurs when the ONTAP® kernel detects an inconsistent inode on disk. Access to this inode is disabled until you repair it by running wافلiron.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

WAFL inconsistent: %s: bad inode (%s fileid: %u snapid: %u) in volume %s%s%s%s.

## Parameters

**func** (STRING): Internal function name.  
**ino\_type** (STRING): Inode type.  
**fileid** (INT): Identifying number for the suspect inode.  
**snapid** (INT): Snapshot(tm) copy identifier (0) - active file system (non-zero) - refers to Snapshot copy.  
**owner** (STRING): Owner of the volume.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

## wافل.inconsistent.ino.tid



**Severity**

ERROR

**Description**

This message occurs when the WAFL® kernel detects an inconsistent inode on a disk with a bad qtree ID.

**Corrective Action**

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. The volume is inconsistent. Run wafiron as soon as possible to clear this inconsistency. For a SnapMirror target, rebuild from the source or resynchronize. If the inode is in a Snapshot copy, delete the Snapshot copy. Then contact technical support.

**Syslog Message**

WAFL inconsistent: bad qtree ID in inode (fileid:%u snapid:%u) in volume %s%. Turning quotas off on the volume.

**Parameters**

**fileid** (INT): Identifying number for the suspect inode

**snapid** (INT): Snapshot(tm) copy identifier (0) - active file system (non-zero) - refers to Snapshot copy

**owner** (STRING): Owner of the volume

**vol** (STRING): Name of the volume

**wafl.inconsistent.snapshot****Severity**

ERROR

**Description**

This event occurs when the WAFL kernel discovers that a snapshot has either an invalid inofile, fsinfo, or snapmap file, or the snapshot vbn stored is invalid.

**Corrective Action**

Delete the snapshot. For a snapmirror target, rebuild from source or resync.

**Syslog Message**

WAFL inconsistent: deleting bad snapshot in volume %s%s%s%s, snapid %d, reason: %s

**Parameters**

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**snapid** (INT): The snapshot identifier

**reason** (STRING): The reason the snapshot is bad

**wafl.inconsistent.stream****Severity**

ERROR

## Description

This event occurs when the WAFL® kernel detects an inconsistent stream information on disk. Access to the stream is disabled until you repair it by running wafllron.

## Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. The volume is inconsistent. Run wafllron as soon as possible to clear this inconsistency; for a snapmirror target, rebuild from source or resync or if the inode is in a snapshot, delete the snapshot. Then contact NetApp technical support.

## Syslog Message

WAFL inconsistent: bad stream (sfileid:%u snapid: %u) for inode (fileid:%u) in volume %s%s%s%s.

## Parameters

**sfileid** (INT): The identifying number for the suspect stream

**snapid** (INT): Snapshot identifier (0) - active file system (non-zero) - refers to snapshot

**fileid** (INT): The identifying number for the inode

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wafll.inconsistent.threshold.reached

### Severity

EMERGENCY

### Description

This event occurs when the number of inconsistencies the WAFL® kernel detected and avoided reached the threshold.

### Corrective Action

The volume is inconsistent. Run wafllron as soon as possible to clear this inconsistency; for a snapmirror target, rebuild from source or resync or if the inode is in a snapshot, delete the snapshot. Then contact NetApp technical support.

### Syslog Message

WAFL inconsistent: The threshold for number of %s inconsistencies was reached. Volume %s %s%s%s%s will be taken offline. Contact NetApp technical support.

### Parameters

**reason** (STRING): Reason : snapshot/inode/block

**type** (STRING): The container type (volume / aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**vendor** (STRING): The vendor name

## wafll.inconsistent.vol

## Severity

ERROR

## Description

This event occurs when the WAFL® kernel discovers an inconsistent volume.

## Corrective Action

The volume is inconsistent. Run wafliron as soon as possible to clear this inconsistency; For a SnapMirror® target, rebuild from source or resynchronize. If the inode is in a Snapshot copy, delete the Snapshot copy. Then contact NetApp technical support.

## Syslog Message

WAFL inconsistent: %s %s%s%s%s is inconsistent. Note: Any new Snapshot copies might contain this inconsistency.

## Parameters

**type** (STRING): The container type (volume / aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**vendor** (STRING): The vendor

## wafl.inconsistent.vol.nomnt

### Severity

ERROR

### Description

This message occurs when the ONTAP® kernel discovers an inconsistent volume and determines that the volume should not be mounted.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

%s %s%s is not mounted.

### Parameters

**type** (STRING): Container type (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

## wafl.ino events

### wafl.ino.upgrade.aborted

### Severity

INFORMATIONAL

**Description**

This message occurs when an inode upgrade for the specified WAFL volume or aggregate is aborted.

**Corrective Action**

(None).

**Syslog Message**

Inode upgrade of volume or aggregate '%s %s' has been aborted.

**Parameters**

**owner** (STRING): Owner of the volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

**wafli.no.upgrade.cantAbort****Severity**

NOTICE

**Description**

This message occurs when an attempt to abort the inode upgrade scan by using the "wafli scan abort" command. The inode upgrade scan cannot be interrupted.

**Corrective Action**

Wait for the inode upgrade to finish or use the "{ vol | aggr } inode-upgrade suspend" command to stop the scan.

**Syslog Message**

User cannot abort the inode upgrade scan. Use the '{ vol | aggr } inode-upgrade suspend' command.

**Parameters**

(None).

**wafli.no.upgrade.completed****Severity**

INFORMATIONAL

**Description**

This message occurs when the inode upgrade of the specified WAFL volume or aggregate is complete.

**Corrective Action**

(None).

**Syslog Message**

An inode upgrade to version %d for volume or aggregate '%s %s' has completed.

**Parameters**

**to\_version** (INT): Version the inode upgraded to.

**owner** (STRING): Owner of the volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

## **waf1.ino.upgrade.disabled**

### **Severity**

ALERT

### **Description**

This message occurs when the inode upgrade is disabled on the node via the "waf1\_v4ino\_upgrade\_disable" bootarg. All aggregates and volumes are still usable, but there might be some features that cannot be enabled on certain volumes until the inode upgrade is completed on those volumes.

### **Corrective Action**

Contact technical support for assistance.

### **Syslog Message**

Inode upgrade is disabled on this node.

### **Parameters**

(None).

## **waf1.ino.upgrade.rem.mdata**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when inode upgrade metadata was removed either by administrator actions or due to waf1iron.

### **Corrective Action**

(None).

### **Syslog Message**

Inode upgrade metadata for volume or aggregate '%s %s' has been removed.

### **Parameters**

**owner** (STRING): Owner of the volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

## **waf1.ino.upgrade.resumed**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an inode upgrade of the specified WAFL volume or aggregate is resumed.

### **Corrective Action**

(None).

## Syslog Message

An inode upgrade for volume or aggregate '%s %s' was resumed.

## Parameters

**owner** (STRING): Owner of the volume or aggregate.  
**volname** (STRING): Name of the volume or aggregate.

## wافل.ino.upgrade.stalled

### Severity

NOTICE

### Description

This message occurs when the administrator suspended an inode upgrade for the specified WAFL volume or aggregate sometime in the past. The volume or aggregate is still usable, but there might be some features that cannot be enabled until the inode upgrade is complete.

### Corrective Action

Use the command 'volume inode-upgrade resume "volume"' or 'storage aggregate inode-upgrade resume "aggregate"' to resume the upgrade.

## Syslog Message

Inode upgrade of volume or aggregate '%s %s' is suspended.

## Parameters

**owner** (STRING): Owner of the volume or aggregate.  
**volname** (STRING): Name of the volume or aggregate.

## wافل.ino.upgrade.started

### Severity

INFORMATIONAL

### Description

This message occurs when the inode upgrade for the specified WAFL volume or aggregate begins.

### Corrective Action

(None).

## Syslog Message

An inode upgrade from version %d for volume or aggregate '%s %s' has begun.

## Parameters

**from\_version** (INT): Current inode version.  
**owner** (STRING): Owner of the volume or aggregate.  
**volname** (STRING): Name of the volume or aggregate.

## wافل.ino.upgrade.susp

## Severity

INFORMATIONAL

## Description

This message occurs when an inode upgrade for the specified WAFL volume or aggregate is suspended by the administrator or because wafiron is running. The volume or aggregate is still usable, but there might be some features that cannot be enabled until the inode upgrade is complete.

## Corrective Action

(None).

## Syslog Message

An inode upgrade for volume or aggregate '%s %s' was suspended. Reason %s. Metadata was %s.

## Parameters

**owner** (STRING): Owner of the volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

**reason** (STRING): Reason the upgrade was suspended.

**metadata\_removed** (STRING): Atring indicating whether metadata was removed as part of the suspension.

## waf.ino.upgrade.susp.nospc

### Severity

ERROR

### Description

This message occurs when an inode upgrade for the specified WAFL volume or aggregate is suspended due to a lack of space. The volume or aggregate is still usable, but there might be some features that cannot be enabled until the inode upgrade is complete.

### Corrective Action

Ensure that the volume or aggregate has at least as many bytes available for use as indicated by 'space\_required'. Simply add the space to the volume and the upgrade is automatically retried.

### Syslog Message

An inode upgrade for volume or aggregate '%s %s' was suspended due to a lack of space (space required: %d bytes).

### Parameters

**owner** (STRING): Owner of the volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

**space\_required** (LONGINT): Estimated amount of disk space required, in bytes, to complete the upgrade.

## waf.inode events

### waf.inode.cache.highLoad

### Severity

INFORMATIONAL

## Description

This message occurs when the in-memory inode cache is under high load and fully utilized. The system will likely experience high cache miss rates and some file system operations might experience delays due to increased disk activity, but that is the extent of the problem.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

(None).

## wافل.inode.fileAttr.fail

### Severity

ERROR

### Description

This message occurs when the system discover that an inode to file attribute translation is unsuccessful due to an unexpected file type. The fh parameters represent the filehandle.

### Corrective Action

The filesystem being accessed is inconsistent. If the filesystem is read-write and if the system is booted in read-only mode then reboot in normal (read-write) mode and run wافلiron as soon as possible to clear this inconsistency. For all other cases contact NetApp technical support.

### Syslog Message

wافل\_inode\_to\_file\_attr: unexpected file type 0%o in %s {x%x %d %d %d %d}.

### Parameters

**type** (INT): The unexpected type  
**ino\_type** (STRING): The inode type  
**fh0** (INT): fh[0]  
**fh1** (INT): fh[1]  
**fh2** (INT): fh[2]  
**fh3** (INT): fh[3]  
**fh4** (INT): fh[4]

## wافل.inode.fill.noDisblRsvFO

### Severity

NOTICE

### Description

This message occurs when the system cannot disable the fill reservation for a file or LUN. This happens when the file or LUN is the destination of an ongoing file operation that is currently in the Allocation\_Map phase.



## Corrective Action

Verify that the indicated file operation has completed the Allocation\_Map phase and rerun the original command.

## Syslog Message

Cannot disable fill reservation for inode %u in volume %s%s%s due to %s job %s-%u.

## Parameters

**ino** (INT): Inode number.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**type** (STRING): Operation type: file copy, file move, LUN copy, or LUN move.

**job\_uuid** (STRING): Job UUID of the file operation.

**file\_index** (INT): File index of the file operation.

# wافل.inofile events

## wافل.inofile.buf.badHeader

### Severity

ERROR

### Description

This event can occur when a snapshot's inofile buffer header has the wrong magic or file block number.

### Corrective Action

Contact customer support.

### Syslog Message

Bad header in snapshot %s inofile in %s %s%s%s%s, snapid %d, fbn %llu, vbn %llu, cp\_count %u was magic 0x%x, fbn %llu, cp\_count %u

### Parameters

**ino\_type** (STRING): The type of inofile

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**snapid** (INT): The snapshot identifier

**fbn** (LONGINT): The file block number

**vbn** (LONGINT): The volume block number

**cp\_count** (INT): The volume CP count

**omagic** (INT): The magic number in the header

**ofbn** (LONGINT): The file block number in the header

**ocp\_count** (INT): The CP count in the header

# wافل.iron events

## **wafli.iron.7mode**

### **Severity**

ERROR

### **Description**

This message occurs when wafliiron is running on a transitioned aggregate in precommit phase and there are 7-Mode volumes present. Wafliiron ignores the 7-Mode volumes.

### **Corrective Action**

Run the "storage transition aggregate start" command to convert the aggregate to a cluster aggregate.

### **Syslog Message**

Transitioned aggregate %s has 7-Mode volumes which cannot be ironed.

### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## **wafli.iron.abort.noConToCloud**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the file system consistency checker (wafliiron) aborts because the connection to the cloud service provider is not available.

### **Corrective Action**

Restore the connection to the cloud service provider, then retry the "aggr wafliiron start <aggr\_name>" command.

### **Syslog Message**

File system consistency checker aborted on aggregate "%s". The connection to the cloud service provider is not available.

### **Parameters**

**aggrname** (STRING): Aggregate name.

## **wafli.iron.abort.offline**

### **Severity**

NOTICE

### **Description**

This message occurs when wafliiron is aborted and the aggregate or volume is taken offline. This will occur if the aggregate or volume is inconsistent at the time of abort or if wafliiron was invoked in optional commit mode. If wafliiron was invoked in optional commit mode, all changes are rejected because they might not be consistent.

### **Corrective Action**

No corrective action is required for this notification event. Wafliiron automatically takes an inconsistent aggregate or volume offline. Wafliiron automatically takes the aggregate offline in optional commit mode.

## Syslog Message

Wafiron was aborted on %s%s '%s'; the aggregate or volume was taken offline.

## Parameters

**volinfo** (STRING): Descriptive information about the volume, such as whether it was inconsistent when taken offline.

**voltype** (STRING): Aggregate or flexible volume.

**aggrname** (STRING): Aggregate or flexible volume name.

## waf.iron.abort.offlineFail

### Severity

NOTICE

### Description

This message occurs when wafiron is aborted and taking an inconsistent aggregate or volume offline was not successful.

### Corrective Action

(Call support).

## Syslog Message

Wafiron abort failed to take offline inconsistent %s '%s': %s.

## Parameters

**voltype** (STRING): Aggregate or flexible volume.

**aggrname** (STRING): Aggregate or flexible volume name.

**errmsg** (STRING): Error message.

## waf.iron.abort.offlineMsg

### Severity

NOTICE

### Description

This message occurs when the wafiron operation is aborted. Taking the aggregate offline was not successful.

### Corrective Action

(None).

## Syslog Message

Wafiron abort on %s '%s': %s.

## Parameters

**voltype** (STRING): Aggregate.

**aggrname** (STRING): Aggregate name.

**errmsg** (STRING): Error message.

## **waf.iron.backup.exists**

### **Severity**

NOTICE

### **Description**

This message occurs when a SnapIron backup image already exists on the volume. Wafiron will continue to run without creating another SnapIron backup image.

### **Corrective Action**

(Call support).

### **Syslog Message**

A SnapIron backup image %s already exists on %s.

### **Parameters**

**vol** (STRING): Volume name.

**name** (STRING): Name of the backup image.

## **waf.iron.backup.expired**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a SnapIron backup image expires and will be deleted.

### **Corrective Action**

(None).

### **Syslog Message**

The SnapIron backup image %s of %s expired.

### **Parameters**

**name** (STRING): Name of the backup image being deleted.

**vol** (STRING): Volume name.

## **waf.iron.backup.expiring**

### **Severity**

NOTICE

### **Description**

This message occurs when a SnapIron backup image is close to expiring.

### **Corrective Action**

Pay attention to the integrity of the data on this volume before the backup image expires and is subsequently deleted by automatic deletion of aggregate Snapshot copies. Contact NetApp technical support.

## Syslog Message

The SnapIron backup image %s of %s will expire in %d hour(s).

## Parameters

**name** (STRING): Name of the backup image.

**vol** (STRING): Volume name.

**hours** (INT): Number of hours remaining from the lifetime of the backup image.

## wافل.iron.backup.impossible

### Severity

NOTICE

### Description

This message occurs when wافلiron encounters a problem when creating a SnapIron backup image. The message explains the reason. Wافلiron will continue to run without the SnapIron backup image.

### Corrective Action

(None).

## Syslog Message

Wافلiron could not create a SnapIron backup image on %s, %s.

## Parameters

**vol** (STRING): Volume name.

**reason** (STRING): Reason for not being able to create a SnapIron backup image.

## wافل.iron.backup.removed

### Severity

NOTICE

### Description

This message occurs when a SnapIron backup image is removed. Wافلiron will continue to run without a SnapIron backup image if it is running.

### Corrective Action

(None).

## Syslog Message

The SnapIron backup image %s of %s was deleted.

## Parameters

**name** (STRING): Name of the backup image being deleted.

**vol** (STRING): Volume name.

## wافل.iron.backup.rst.failed

### Severity

ALERT

## Description

The message occurs when the system cannot restore the aggregate from the wafiron backup image. The active file system will be loaded.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

Failed to restore %s %s%s%s%s from its wafiron backup image (ID %d). The active file system will be loaded.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**id** (INT): Backup image ID.

**reason** (STRING): Reason for the failure.

## wafiron.backup.rst.noExist

### Severity

NOTICE

### Description

The message occurs when the system cannot restore the aggregate from the wafiron backup image because the image does not exist. The active file system will be loaded.

### Corrective Action

(None).

### Syslog Message

Could not restore %s %s%s%s%s from its wafiron backup image because the image does not exist. The active file system will be loaded.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wafiron.backup.rst.success

### Severity

INFORMATIONAL

### Description

The message occurs when the system restores the aggregate from its wafiron backup image.

## Corrective Action

(None).

## Syslog Message

Restored %s %s%s%s%s from its wafiron backup image (ID %d).

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**id** (INT): Backup image ID.

## waf.iron.backup.snaprsv.inc

### Severity

INFORMATIONAL

### Description

This message occurs when wafiron tries to increase the Snapshot(tm) reserve of the aggregate temporarily to allow keeping the Snaplron backup image longer. The message shows whether the attempt was successful.

### Corrective Action

(None).

### Syslog Message

The attempt to increase Snapshot copy reserve of %s temporarily from %d percent by %d %s.

### Parameters

**vol** (STRING): Volume name.

**rsv** (INT): Old value of the aggregate Snapshot reserve, in percentage.

**delta** (INT): Change in the new value of the aggregate Snapshot copy reserve, in percentage. If the Snapshot copy reserve was successfully increased, it will be returned to its original value automatically after the Snaplron backup image expires, the new snapshot reserve space becomes full again, or the snapshot is deleted. In the latter case, the Snaplron backup image will be deleted by automatic deletion of aggregate Snapshot copies.

**result** (STRING): Whether the attempt was successful.

## waf.iron.backup.snaprsv.rst

### Deprecated

This is deprecated as success doesn't require any action failure can never happen.

### Severity

NOTICE

### Description

This message occurs when the Snapshot(tm) reserve of the aggregate, which was temporarily increased by wafiron, is restored to its original value.

## Corrective Action

If wafiron fails to return the volume's Snapshot copy reserve to the original value, contact NetApp technical support.

## Syslog Message

The reset of Snapshot reserve value of %s by %d from %d was %s.

## Parameters

**vol** (STRING): Volume name.

**val** (INT): Current value of the aggregate Snapshot copy reserve, in percentage.

**delta** (INT): Change in the value of the aggregate Snapshot reserve, in percentage.

**result** (STRING): Whether the attempt was successful.

## waf.iron.backup.success

### Severity

INFORMATIONAL

### Description

This message occurs when wafiron creates a SnapIron backup image.

### Corrective Action

(None).

### Syslog Message

Wafiron created SnapIron backup image %s on volume %s, ID %d.

### Parameters

**name** (STRING): Backup image name.

**vol** (STRING): Volume name.

**id** (INT): Backup image ID.

## waf.iron.backup.useimage

### Severity

INFORMATIONAL

### Description

This message occurs when wafiron is using a SnapIron backup image to check the file system.

### Corrective Action

(None).

### Syslog Message

Wafiron is using SnapIron backup image of %s (ID %u).

### Parameters

**vol** (STRING): Volume name.

**id** (INT): ID of SnapIron backup image.



## **waf.iron.backup.verify.data**

### **Severity**

NOTICE

### **Description**

This message occurs when wafiron creates a SnapIron backup image and you need to verify whether user data is stored in the aggregate or volume as soon as possible.

### **Corrective Action**

Validate the user data stored in the aggregate or volume as soon as possible, before the SnapIron backup image is deleted. Contact NetApp technical support.

### **Syslog Message**

Verify user data of %s.

### **Parameters**

**vol** (STRING): Volume name.

## **waf.iron.badfsid**

### **Severity**

ALERT

### **Description**

This message occurs when a flexible volume has an invalid file system ID(FSID). Wafiron skips this volume.

### **Corrective Action**

Contact NetApp technical support.

### **Syslog Message**

Flexible volume %s%s%s on aggregate %s has an invalid FSID (%08x).

### **Parameters**

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**aggrname** (STRING): Aggregate name.

**fsid** (INT): File system identifier.

## **waf.iron.completion.times**

### **Severity**

INFORMATIONAL

### **Description**

This event provides the time taken to complete the different phases of Wafiron.

### **Corrective Action**

(None).

## Syslog Message

%s phase of %s %s took %s.

## Parameters

**p** (STRING): Name of the iron phase.  
**ty** (STRING): Type of the volume.  
**v** (STRING): Name of the volume.  
**t** (STRING): Time taken for the phase.

## wافل.iron.cont.corrupt

### Severity

ALERT

### Description

This message occurs when a volume granular wافلiron operation running on an aggregate finds corruption in the container file of a volume on which wافلiron is not being run.

### Corrective Action

Contact Contact NetApp technical support..

## Syslog Message

Volume granular wافلiron is in progress on aggregate %s. Corruption was found in volume %s container file with file id %d in which wافلiron might have not run.

## Parameters

**aggr\_name** (STRING): Name of the aggregate on which wافلiron is running.  
**vol\_name** (STRING): Name of the volume in which corruption was found in the container file.  
**container\_file\_id** (INT): File id of container file of the volume in which corruption was found.

## wافل.iron.current.phase

### Severity

INFORMATIONAL

### Description

This message specifies the phase name that wافلiron is currently working on.

### Corrective Action

(None).

## Syslog Message

wافلiron is active on %s %s : %s.

## Parameters

**ty** (STRING): Type of the volume.  
**v** (STRING): Name of the volume.  
**phase** (STRING): Name of the current wافلiron phase.

## **wافل.iron.dp.mirror.broken**

### **Severity**

ERROR

### **Description**

This message occurs when a Data Protection mirrored volume is broken from its SnapMirror® relationship as a result of starting wافلiron with the '-include-mirrors' option. When this happens, the Volume Location Database (VLDB), which is managed by the management host, must also be updated manually after wافلiron finishes to avoid any discrepancy between the VLDB and Dblade.

### **Corrective Action**

Run the 'debug vreport' command to resolve the inconsistency in the VLDB. Step 1) run 'debug vreport show' to show the object name. Step 2) run 'debug vreport fix -type volume -object object-name'.

### **Syslog Message**

WافلIron broke the SnapMirror® relationship for volume '%s%s%s'. Use the 'debug vreport' command to resolve inconsistencies before reestablishing the SnapMirror relationship.

### **Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## **wافل.iron.grow.fs.failed**

### **Severity**

ERROR

### **Description**

This message occurs when WAFL® cannot grow a volume or aggregate because wافلiron is running and the root volume does not have enough free space to grow the iron status files.

### **Corrective Action**

Check the free space in the root volume. Abort wافلiron or grow the root volume to accommodate the grown status files

### **Syslog Message**

WAFL failed to grow %s %s due to a failure in growing wافلiron status files in the root volume.

### **Parameters**

**voltype** (STRING): Volume type: volume or aggregate.

**volname** (STRING): Name of the volume or aggregate.

## **wافل.iron.i2p.dir.convert**

### **Severity**

INFORMATIONAL

## Description

This message occurs to inform you why wafiron is starting the inode-to-parent (i2p) initialization scan. The i2p scan updates the information in the volume necessary to perform i2p path name translations. The reason i2p information needs updating is that the volume contains converted directories, and the entries in the converted directories need their i2p information updated. This message does not indicate anything abnormal or problematic.

## Corrective Action

(None).

## Syslog Message

Found converted directories in %s %s%s%s%s that require their entry's inode-to-parent (i2p) information updated. Starting the i2p initialization scanner to update the entries in the converted directories. This is normal.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf.iron.incompleteVol

### Severity

ALERT

### Description

This wafiron event indicates that a volume is partially created and might be in an inconsistent state. Wafiron is marking the volume as unrecoverable. For instance, if filer panics before snap restore completes when a clone volume is created it may leave the volume partially created.

### Corrective Action

If the volume is known not to contain any valuable data, it can be safely destroyed and recreated.

### Syslog Message

Flexible volume %s%s on aggregate %s is partially created. Marking it unrecoverable. You may destroy this volume and recreate it.

### Parameters

**volname** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**aggrname** (STRING): Aggregate name

## waf.iron.invalid.vvol

### Severity

ERROR

## Description

This wafiron message indicates that a volume is invalid. Wafiron cannot run on invalid volumes. A volume can become invalid as a result of an aborted 'vol copy' operation or SnapMirror® initial transfer. To bring the volume back online, you must re-create it through a completed 'vol copy' operation or SnapMirror transfer.

## Corrective Action

If you know that the volume does not contain any valuable data, you can safely destroy and re-created it.

## Syslog Message

Flexible volume %s%s%s on aggregate %s is invalid. Wafiron cannot run on invalid volumes.

## Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**aggrname** (STRING): Aggregate name.

## wafl.iron.lost.blocks

### Severity

INFORMATIONAL

### Description

This message indicates that wafiron is running on a particular inode.

### Corrective Action

(None).

### Syslog Message

Wafiron on %s %s%s%s is verifying %lld %s inode now.

### Parameters

**voltype** (STRING): Volume type.

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**inode\_number** (LONGINT): Number of inode buftrees not verified.

**ino\_type** (STRING): The inode type.

## wafl.iron.manualUpgradeRequired

### Severity

ERROR

### Description

This event is issued when wafiron can't run on the volume because it's not upgraded yet. This happens when system was upgraded with an offlined or restricted volume (possibly due to corruption) and wafiron is started after the upgrade with the volume offlined or restricted.

### Corrective Action

Manually online the volume and then restart wafiron. Or fix any corruption on an inconsistent volume and bring it online before a kernel upgrade.

## Syslog Message

Wafiron cannot run on volume '%s' that is in an older format. Manually online the volume and retry wafiron.

## Parameters

**volname** (STRING): Volume name.

## waf.iron.mount.inconsistency

### Severity

ERROR

### Description

This message occurs when wafiron finds some WAFL® inconsistency issues on an aggregate during mount.

### Corrective Action

Run the 'storage aggregate wafiron review <aggregate-name>' command to review the issues that wafiron found. If you need assistance resolving any of the issues, contact NetApp technical support.

## Syslog Message

wafiron on aggregate %s found some WAFL inconsistency issues during mount.

## Parameters

**aggrname** (STRING): Aggregate name.

## waf.iron.mount.inconsistent

### Severity

ALERT

### Description

This wafiron message indicates that the FlexVol® volume cannot be mounted; the volume is marked as inconsistent.

### Corrective Action

Re-run wafiron with the '-prev\_cp' option set to restore the FlexVol volume to the closest Consistency Point(CP)/Snapshot(tm) copy.

## Syslog Message

wafiron could not mount volume %s due to %s. Volume is marked as inconsistent.

## Parameters

**volname** (STRING): Volume name.

**reason** (STRING): Reason for the failure.

## waf.iron.mount.times

### Severity

INFORMATIONAL

## Description

This message prints the mount times taken by the different phases of aggregate mount during wafiron. This message occurs after wafiron completion or through a wafiron status message.

## Corrective Action

(None).

## Syslog Message

%s mount phase of %s %s took %s.

## Parameters

**phase** (STRING): Name of the mount phase.  
**voltype** (STRING): Type of the volume (aggragate/flexvol).  
**volname** (STRING): Name of the volume or aggregate.  
**time** (STRING): Time taken for the phase.

## waf.iron.mounted.volume

### Severity

INFORMATIONAL

### Description

This message is used to print on the console that the volume is mounted during wafiron.

### Corrective Action

(None).

### Syslog Message

Wafiron mounted volume %s.

### Parameters

**vol** (STRING): Name of the volume.

## waf.iron.mw.ioc.not.supported

### Severity

ERROR

### Description

This event is issued when IOC is issued on metawafI aggregate.

### Corrective Action

(Call support).

### Syslog Message

Iron optional commit is not supported on metawafI aggregate '%s'. Please run wafiron.

### Parameters

**volname** (STRING): Volume name.

## **wagl.iron.nvsave.disable.selection**

### **Severity**

INFORMATIONAL

### **Description**

This message is used to print on the console that the wafiron cannot run in volume select mode on the aggr as the aggregate has nvsave file records.

### **Corrective Action**

(None).

### **Syslog Message**

The aggr %s has nvlog records in the system. Wafiron cannot run in volume selection mode. Volumes will be brought online in random order.

### **Parameters**

**aggr** (STRING): Name of the aggregate.

## **wagl.iron.oc.abort.bad\_blk**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the system aborts wafiron in optional commit mode on the volume due to a bad block in the change log and lost write context status files.

### **Corrective Action**

Refer to the Recovery Decision Tree.

### **Syslog Message**

Aborting wafiron in optional commit mode because a bad block was found in volume '%s'.

### **Parameters**

**volname** (STRING): Volume name.

## **wagl.iron.oc.abort.clog\_full**

### **Severity**

EMERGENCY

### **Description**

This message occurs when the system aborts an optional commit wafiron operation on a volume because the change log file was full, due to extensive filesystem consistency changes.

### **Corrective Action**

Contact NetApp technical support.

### **Syslog Message**

Aborting wafiron in optional commit mode on aggregate '%s'. The change log is full.



**Parameters**

**volname** (STRING): Volume name.

**wafli.iron.oc.committedChangeLog****Severity**

NOTICE

**Description**

This wafliiron event indicates that optional commit change log files have been committed to disk.

**Corrective Action**

This happens when the user explicitly commits the changes to disk.

**Syslog Message**

Committed change log files for aggregate %s.

**Parameters**

**volname** (STRING): Volume name

**wafli.iron.oc.deletedChangeLog****Severity**

EMERGENCY

**Description**

This wafliiron event indicates that optional commit change log files have been deleted.

**Corrective Action**

This happens when the change log files have been written to disk, or have been rejected.

**Syslog Message**

Deleted change log files for aggregate %s.

**Parameters**

**volname** (STRING): Volume name

**wafli.iron.oc.errorCommitLog****Severity**

EMERGENCY

**Description**

This wafliiron event indicates that optional commit change log files' commit process has encountered an error.

**Corrective Action**

Corrective action may include offlining the aggregate and running wafliiron in optional commit mode again.

**Syslog Message**

Error committing change log files for aggregate %s.

**Parameters**

**aggrname** (STRING): Aggregate name

**wafli.iron.oc.noCorruption****Severity**

NOTICE

**Description**

This wafliiron event indicates that optional commit did not find any corruption in the aggregate that it is being run on.

**Corrective Action**

No corrective action is required for this notification event. Wafliiron will automatically online the aggregate.

**Syslog Message**

Wafliiron did not find any corruption on aggregate %s. The aggregate and its volumes will be offlined automatically.

**Parameters**

**aggrname** (STRING): Aggregate name

**wafli.iron.oc.rejectChangeLog****Severity**

NOTICE

**Description**

This wafliiron event indicates that optional commit change log files have been rejected (not committed).

**Corrective Action**

No corrective action required for this notification event.

**Syslog Message**

Rejecting change log files for aggregate %s.

**Parameters**

**aggrname** (STRING): Aggregate name

**wafli.iron.oc.root.committedChanges****Severity**

NOTICE

**Description**

This wafliiron event indicates that optional commit changes have been committed to disks of the root aggregate/tradvol.

**Corrective Action**

No corrective action required.

**Syslog Message**

Committed changes for %s '%s'. The system will reboot now.

**Parameters**

**voltype** (STRING): Volume type - aggregate or tradvol

**volname** (STRING): Volume name

**wafli.iron.oc.root.lowMemory****Severity**

EMERGENCY

**Description**

Wafliiron Optional Commit on the Root has exhausted available memory and cannot continue. Options will be presented to the administrator to either abort the changes or write to disk.

**Corrective Action**

Changes will have to be committed to disk for wafliiron to continue or will have to be rejected.

**Syslog Message**

Wafliiron Optional Commit on '%s' has run out of available memory and cannot continue. Please either commit to continue Wafliiron by writing changes to disk or reject the changes (which will reboot the system).

**Parameters**

**volname** (STRING): Volume name

**wafli.iron.oc.root.noCorruption****Severity**

NOTICE

**Description**

This wafliiron message indicates that the optional commit operation does not find any corruption in the root aggregate that it is being run on.

**Corrective Action**

No corrective action is required.

**Syslog Message**

Wafliiron did not find any corruption on aggregate %s. The system will reboot now.

**Parameters**

**aggname** (STRING): Root aggregate name.

**wafli.iron.oc.root.rejectedChanges**

**Severity**

NOTICE

**Description**

This wafiron event indicates that optional commit changes have been rejected.

**Corrective Action**

No corrective action required.

**Syslog Message**

Rejected changes for %s '%s'. The system will reboot now.

**Parameters**

**voltype** (STRING): Volume type - aggregate or tradvol

**volname** (STRING): Volume name

**waf.iron.offline.destroying****Severity**

INFORMATIONAL

**Description**

This message occurs when a volume is being destroyed and cannot be ironed/online.

**Corrective Action**

(None).

**Syslog Message**

Volume '%s%s%s' is being destroyed and was taken offline.

**Parameters**

**volname** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**waf.iron.optcommit.upgradeScan****Severity**

ERROR

**Description**

This message occurs when wafiron optional commit is started in a volume that has an upgrade still in progress.

**Corrective Action**

Wait for the upgrade scanner to finish for this volume, then restart wafiron.

**Syslog Message**

Wafiron optional commit cannot run in volume '%s' that still has an upgrade in progress. Run the advanced CLI command 'waf scan status' to check whether an upgrade is still in progress.

## Parameters

**volname** (STRING): Volume name.

## **wافل.iron.outofspace**

### Severity

ERROR

### Description

This message occurs when an aggregate is out of space. Wafiron omits the volumes in this aggregate when managing the aggregate.

### Corrective Action

Run the "storage transition commit start" command to commit the aggregate transition and delete the pinned Snapshot® copy.

### Syslog Message

Volumes in the aggregate %s are kept offline since aggregate is out of space. Wafiron omits volumes in this aggregate.

## Parameters

**aggrname** (STRING): Name of the aggregate.

## **wافل.iron.parallel.mount.notify**

### Severity

NOTICE

### Description

This message indicates that wafiron parallel mount is proceeding for the flexible volumes in an aggregate. All sparse volumes are mounted in parallel first, followed by non-clone flexible volumes in parallel. Each clone is mounted following its parent volume.

### Corrective Action

(None).

### Syslog Message

Wafiron parallel mount enabled. Wafiron is mounting flexible volumes in parallel in the aggregate %s.

## Parameters

**aggr** (STRING): Aggregate name

## **wافل.iron.parent.volume.selected**

### Severity

INFORMATIONAL

### Description

This message occurs when the parent of a clone volume is added to the list of volumes wafiron will mount first because the clone volume is on the list.

## Corrective Action

(None).

## Syslog Message

Volume %s in aggregate %s was added to the list of volumes that wafiron will mount first because its clone volume %s is present in this list.

## Parameters

**parent\_volume** (STRING): Name of the parent volume.

**aggr** (STRING): Name of the aggregate.

**clone\_volume** (STRING): Name of the clone volume.

## waf.iron.pending.delete.vvol

### Severity

ERROR

### Description

This wafiron event indicates that a volume is awaiting deletion and cannot be ironed/online. Its deletion also cannot be canceled while iron is running, so its retention time may need to be extended in order to keep the volume.

### Corrective Action

If you do not want the volume to be deleted while iron is running, check the retention time and consider extending it.

### Syslog Message

Flexible volume %s%s%s on aggregate %s is awaiting deletion. Wafiron cannot run on such volumes.

### Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**aggrname** (STRING): Aggregate name.

## waf.iron.ral.exception

### Severity

ERROR

### Description

This message occurs when wafiron is loading an inode that has a remote entry, but the corresponding remote volume is not reachable.

### Corrective Action

Inspect the status of the FlexGroup remote volume, which might be down due to network errors, cluster node failures, or corruption in one or more of its constituent aggregates. Also check the health and connectivity of all nodes that are hosting these aggregates.

### Syslog Message

Wafiron could not load inode "%d" on volume "%s" because the remote volume is not accessible.

## Parameters

**inode** (INT): Inode that could not be loaded.  
**volume** (STRING): Volume being inspected by wafliiron.

## wafliiron.ral.inode.loaded

### Severity

NOTICE

### Description

This message occurs when wafliiron succeeds in loading an inode in a FlexGroup origin volume that was not reachable earlier.

### Corrective Action

NONE.

### Syslog Message

After several retries, wafliiron has processed inode "%d" in volume "%s".

## Parameters

**inode** (INT): Inode loaded after retries.  
**volume** (STRING): Volume being inspected by wafliiron.

## wafliiron.readonly

### Severity

NOTICE

### Description

This message occurs during wafliiron volume initialization, indicating that read-only volumes might not be checked.

### Corrective Action

Use either the "-f" or the "-include-mirrors" option to run wafliiron on SnapMirror® destination volumes. For more information or assistance, contact NetApp technical support.

### Syslog Message

%s%s%s: Readonly volumes require additional command options to be checked using wafliiron

## Parameters

**owner** (STRING): Owner of the volume.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## wafliiron.readonlyQtrees

### Severity

NOTICE

## Description

This message occurs during wafiron volume initialization, indicating that volumes with read-only qtrees might not be checked.

## Corrective Action

Use either the "-f" or the "-include-mirrors" option to run wafiron on SnapMirror® destination volumes.

## Syslog Message

%s%s%s: Volumes containing readonly qtrees require additional command options to be checked using wafiron

## Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## waf.iron.rebuild.dir.volful

### Severity

ERROR

### Description

This message occurs while running wafiron, when there might not be enough free space on the volume to rebuild a damaged directory.

### Corrective Action

Increase the size of the volume or free space on the aggregate that contains the volume.

### Syslog Message

Volume %s might run out of space while wafiron is rebuilding the damaged directory %d. The required space is %d MB.

### Parameters

**vol** (STRING): Name of the volume.

**fileid** (INT): Identifying number for the damaged directory.

**spaceNeeded** (INT): Amount of spaced needed, in megabytes.

## waf.iron.salvaged.rootdir

### Severity

INFORMATIONAL

### Description

This message occurs when wafiron detects a corrupted rootdir in an index-dir-enabled volume and re-creates it with all valid entries in the rootdir. The old corrupted rootdir has been converted into a regular file and linked to the /lost+found directory.

### Corrective Action

Remove the file from the /lost+found directory.



## Syslog Message

Rootdir of volume %s has been corrected and contains all valid names. The contents of the old rootdir have been copied to file i#%d.\* in the /lost+found directory.

## Parameters

**volume** (STRING): Name of the volume.

**fileid** (INT): File ID of the regular file that contains the old rootdir contents as is.

## wافل.iron.scan.skipped

### Severity

NOTICE

### Description

This event occurs when non-specified or read only volumes are not processed during a wافلiron run on their containing aggregate.

### Corrective Action

(None).

## Syslog Message

Aggregate %s, wافلiron completed - Read only, offline, old version or non-specified volume(s) %s not checked.

## Parameters

**aggr** (STRING): The aggr name

**vvols** (STRING): Flexible volumes not processed

## wافل.iron.scan.stats

### Severity

INFORMATIONAL

### Description

This message prints the wافلiron scan details about various data objects: for example, files, directories, blocks.

### Corrective Action

(None).

## Syslog Message

%llu %s done.

## Parameters

**count** (LONGINT): Number of data objects scanned so far.

**object** (STRING): Name of the data object.

## wافل.iron.scanning.status

**Severity**

INFORMATIONAL

**Description**

This message prints wafiron scanning status.

**Corrective Action**

(None).

**Syslog Message**

Scanning (%d%% done).

**Parameters**

**percentage** (INT): Scanning percentage.

**waf.iron.slow****Severity**

ERROR

**Description**

This message occurs when a "wafiron start" or a "wafiron show" command is issued and the aggregate that hosts wafiron status files for the aggregate on which wafiron is running contains only one data disk.

**Corrective Action**

1. Use the "-status-files-vserver" and "-status-files-volume" parameters of the "aggr wafiron start" command to host status files on a non-root aggregate volume, that has an adequate number of data disks that can accommodate the extra workload. 2. Add more disks to the root aggregate using "aggr add" command.

**Syslog Message**

Aggregate %s: wafiron performance will be significantly impacted because aggregate %s contains only one data disk. Running wafiron with status files hosted on a non-root aggregate volume that has an adequate number of data disks, or adding more disks to the root aggregate is recommended.

**Parameters**

**aggr\_name** (STRING): Name of the aggregate on which wafiron is running.

**status\_vol\_aggr\_name** (STRING): Name of the aggregate containing the volume hosting the wafiron status files.

**waf.iron.start****Severity**

INFORMATIONAL

**Description**

This message occurs when the system starts the wafiron process on a volume.

**Corrective Action**

(None).

## Syslog Message

Starting wafliiron on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**optcommit** (STRING): Whether wafliiron is running in optional-commit mode. If this mode is used, the decision to commit changes is made by the user after wafliiron is complete.

**prev\_cp** (STRING): Whether the prev\_cp feature of wafliiron is being used and, if so, which consistency point the user chose at the prompt.

**backup** (STRING): Whether the backup feature of wafliiron is being used and, if so, what option {no|yes|force|use} is selected. Only applicable to aggregates.

## wafliiron.unrecoverable

### Severity

ALERT

### Description

This message occurs when a volume is unrecoverable and cannot be ironed or brought online.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Flexible volume %s%s on aggregate %s is unrecoverable.

## Parameters

**volname** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**aggrname** (STRING): Aggregate name.

## wafliiron.upgradeRequired

### Severity

NOTICE

### Description

This message occurs during volume or aggregate initialization if wafliiron is enabled and does not support the volume or aggregate file system version. The message indicates that a volume or aggregate upgrade is required and is being attempted without wafliiron checking. A wafliiron check of the volume or aggregate might be possible after a successful upgrade.

### Corrective Action

(None).

## Syslog Message

%s '%s%s%s%s' has not been upgraded. Upgrade will proceed without wafiron. The WAFL version on the %s is '%d'; the current WAFL version is '%d'.

## Parameters

**Type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume or aggregate.

**vol** (STRING): Name of the volume or aggregate.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume or aggregate in cases where the volume or aggregate name itself is insufficient.

**type** (STRING): Type of object (volume or aggregate).

**vol\_wafl\_version** (INT): WAFL® version on the volume or aggregate.

**current\_wafl\_version** (INT): Current WAFL version.

## waf.iron.volumes.selected.not.exist

### Severity

INFORMATIONAL

### Description

This message is used to print on the console that the volumes selected for ironing on the aggregate do not exist in the aggregate.

### Corrective Action

(None).

## Syslog Message

Some of the volumes selected for ironing do not exist in the aggregate %s.

## Parameters

**aggr** (STRING): Name of the aggregate.

## waf.latest.events

## waf.latest.physical.snapid.info

### Severity

INFORMATIONAL

### Description

This message occurs when WAFL® detects an inconsistent physical Snapshot(tm) copy ID for the latest Snapshot copy on a volume or an aggregate.

### Corrective Action

(None).

## Syslog Message

Corrected the latest\_snapid for %s %s from %d to %d.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**volume** (STRING): Volume or aggregate name.  
**invalid\_latest\_snapid** (INT): Invalid latest Snapshot copy ID stored on disk.  
**corrected\_latest\_snapid** (INT): The corrected latest Snapshot copy ID.

# wafl.Imgr events

## wafl.Imgr.meta.mv

### Severity

NOTICE

### Description

The appliance file lock state for a volume could not be moved during a file system check operation (wafiron).

### Corrective Action

(None).

### Syslog Message

Move of file lock manager state for volume fsid %08x (uuid %s) for file system check received error %s

## Parameters

**fsid** (INT): File system identifier of the volume.  
**uuid** (STRING): Uuid of the volume.  
**err** (STRING): Error message.

## wafl.Imgr.meta.rm

### Severity

NOTICE

### Description

The appliance file lock state for a volume could not be removed.

### Corrective Action

(None).

### Syslog Message

Cannot remove lock manager host metafile %s in recycle bin: error %s

## Parameters

**path** (STRING): Metadata file path.  
**err** (STRING): Error message.

# wafl.lookup events

## wافل.lookup.smDirErr

### Severity

ERROR

### Description

This message occurs when the system encounters an unexpected deleted directory entry in a snapmirrored qtree. This message can occur when accessing a directory in a snapmirrored qtree.

### Corrective Action

This condition is transient. Wait for any in-progress SnapMirror(tm) operations to finish and try the operation again.

### Syslog Message

wافل\_nfs\_lookup: unexpected deleted directory entry in snapmirrored qtree %d: {x%x %d %d %d %d} <%. \*s> in {x%x %d %d %d %d}.

### Parameters

**qtree** (INT): Qtree identifier.  
**fh0** (INT): Directory file handle's flags.  
**fh1** (INT): Directory file handle's snapid.  
**fh2** (INT): Directory file handle's fileid.  
**fh3** (INT): Directory file handle's generation.  
**fh4** (INT): Directory file handle's FSID.  
**name** (STRINGARRAY): Name of the directory.  
**dfh0** (INT): Requested file handle's flags.  
**dfh1** (INT): Requested file handle's snapid.  
**dfh2** (INT): Requested file handle's fileid.  
**dfh3** (INT): Requested file handle's generation.  
**dfh4** (INT): Requested file handle's FSID.

## wافل.make events

### wافل.make.container.dense.failed

### Severity

NOTICE

### Description

This message occurs when container cant be made dense

### Corrective Action

(None).

### Syslog Message

Container cannot be made dense for %s "%s" due to "%s".

### Parameters

**volume\_identifier** (STRING): Volume identifier.  
**volume** (STRING): Volume name.  
**message** (STRING): Reason for failure to make container dense

# wafl.maxdirsize events

## wafl.maxdirsize.boot.notice

### Severity

NOTICE

### Description

This event occurs for each volume that has a maxdirsize greater than the default during boot.

### Corrective Action

Use (privilege:advanced) option "volume modify -volume vol\_name -maxdir-size new\_value" to set the volume's maxdirsize to the default maxdirsize. However, doing so could impact system performance. If you need to increase the maximum directory size, work with technical support.

### Syslog Message

%s%s%s%s: This volume's maxdirsize (%uKB) is higher than the default (%uKB). There may be a performance penalty when doing operations on large directories.

### Parameters

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name alone is insufficient.

**current\_maxdirsize** (INT): Default maximum directory size value.

**default\_maxdirsize** (INT): Default maximum directory size value.

# wafl.mcc events

## wafl.mcc.config.notFound

### Deprecated

Deprecated as of version 9.4 as this debug EMS not supported now.

### Severity

ERROR

### Description

This message occurs when WAFL® does not find the MetroCluster configuration information on this node.

### Corrective Action

Run the 'metrocluster check' command to verify the MetroCluster configuration.

### Syslog Message

(None).

### Parameters

(None).

## **waf1.mcc.exceed.volcnt.limit**

### **Severity**

ERROR

### **Description**

This message occurs when the MetroCluster® volume count limit for this node is reached. No additional volumes can be created on this node unless some existing volumes are moved or destroyed.

### **Corrective Action**

Either destroy some existing volumes or perform a 'volume move' operation to move some existing volumes to a storage aggregate located on a different node to make room for the new ones. Alternatively, you can also create new volumes on a different node to avoid moving or deleting the existing volumes.

### **Syslog Message**

MCC volume count limit of %d reached for this node.

### **Parameters**

**max\_volcnt** (INT): Volume count limit for this node.

## **waf1.mcc.nso.left.behind**

### **Severity**

ERROR

### **Description**

This message occurs when the system aborts a negotiated switchover because an online left-behind aggregate was found on this node.

### **Corrective Action**

Check whether there exists any online left-behind aggregates on this node from a previous switchover or takeover operation. If such an aggregate exists, then perform the corrective actions specified in the earlier EMS messages for the left-behind aggregate to return it to its original owner, and then retry the operation. You can use the "override-vetoes" option to proceed without performing the corrective action.

### **Syslog Message**

One or more online left-behind aggregates exist. Negotiated switchover is canceled.

### **Parameters**

(None).

## **waf1.mcc.nso.veto.cmode.vols**

### **Severity**

ERROR

### **Description**

This message occurs when the system aborts negotiated switchover because a cluster volume was found on an aggregate with Controller Failover (CFO) HA policy.



### Corrective Action

Retry the negotiated switchover operation after the volume is deleted or moved to an aggregate with SFO HA policy. The override-vetoes option may be used to proceed without deleting or moving the volume.

### Syslog Message

%s %s%s%s%s has a CFO HA policy and contains one or more cluster volumes. Negotiated switchover canceled.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wافل.mcc.sb.left.behind

### Severity

ERROR

### Description

This message occurs when WAFL® leaves behind an aggregate during MetroCluster(tm) switchback because this aggregate does not meet all the WAFL veto criteria for switchback.

### Corrective Action

Correct the issue that caused the aggregate to be left behind, based on the specific reason identified in an earlier EMS message. After correcting the issue, run the 'metrocluster switchback' command to return the aggregate to its original owner.

### Syslog Message

Aggregate %s was left behind during switchback.

### Parameters

**vol** (STRING): Name of the aggregate that was left behind.

## wافل.mcc.sb.nso.veto.create

### Severity

ERROR

### Description

This message occurs when the system aborts a switchback or negotiated switchover because a volume create operation is running.

### Corrective Action

Retry the switchback or negotiated switchover operation after the volume create operation is complete.

### Syslog Message

%s %s%s%s%s is being created. Switchback or negotiated switchover canceled.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.mcc.sb.nso.veto.fence

### Severity

ERROR

### Description

This message occurs when the system aborts a switchback or negotiated switchover operation because a volume has a fence set.

### Corrective Action

A volume in the aggregate is currently fenced for group operations. Retry the switchback or negotiated switchover operation after any group SnapRestore® or mirror transfers are complete. If the problem persists, contact technical support.

### Syslog Message

%s %s%s%s%s has a fence set. Switchback or negotiated switchover canceled.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.mcc.sb.nso.veto.iron

### Severity

NOTICE

### Description

This message occurs when the system aborts a switchback or negotiated switchover because wafiron is running.

### Corrective Action

Retry the switchback or negotiated switchover operation after wafiron is finished.

### Syslog Message

%s %s%s%s%s is running wafiron. Switchback or negotiated switchover canceled.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## **waf1.mcc.sb.nso.veto.lock**

### **Severity**

ERROR

### **Description**

This message occurs when WAFL® vetoes a switchback or negotiated switchover operation because the system cannot obtain an aggregate level lock.

### **Corrective Action**

Retry the switchback or negotiated switchover operation after waiting for other operations on the aggregate to finish.

### **Syslog Message**

Switchback or negotiated switchover vetoed because the system could not obtain a lock on aggregate %s.

### **Parameters**

**vol** (STRING): Aggregate name.

## **waf1.mcc.sb.nso.veto.nvsave**

### **Severity**

ERROR

### **Description**

This message occurs when WAFL® vetoes a switchback or negotiated switchover operation because there are NVSave files for an aggregate or a flexible volume that must be replayed before the aggregate or flexible volume can be switched back.

### **Corrective Action**

Retry the switchback or negotiated switchover operation after replaying NVSave files (by bringing the aggregate or flexible volume online) or deleting them from the root volume.

### **Syslog Message**

Switchback or negotiated switchover vetoed due to presence of NVSave files for '%s' (FSID %d, UUID %s).

### **Parameters**

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

**fsid** (INT): File system identifier.

**UUID** (STRING): UUID of the Volume.

## **waf1.mcc.sb.nso.veto.revert**

### **Severity**

ERROR

### **Description**

This message occurs when the system aborts a switchback or negotiated switchover because a volume revert operation is running.

### **Corrective Action**

Retry the switchback or negotiated switchover operation after the volume revert operation is complete.

### **Syslog Message**

%s %s%s%s%s is being reverted. Switchback or negotiated switchover canceled.

### **Parameters**

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## **wafl.mcc.so.replay.disabled**

### **Severity**

NOTICE

### **Description**

This message occurs during MetroCluster(tm) switchover processing on a simulator. WAFL® replay is disabled on simulators.

### **Corrective Action**

(None).

### **Syslog Message**

MetroCluster switchover: WAFL replay disabled on simulators.

### **Parameters**

(None).

## **wafl.mcc.so.start**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when WAFL starts switchover operation.

### **Corrective Action**

(None).

### **Syslog Message**

WAFL starts switchover operation (force-nvfail-all=%s).

### **Parameters**

**force\_nvfail\_all** (STRING): Switchover command option.

## **wafl.mcc.so.vols.mounted**

**Severity**

INFORMATIONAL

**Description**

This message occurs when the volumes are mounted during switchover.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

**waf1.mcc.subphase.times****Severity**

INFORMATIONAL

**Description**

Issued upon completion of the WAFL portion of a MetroCluster operation with sub-phase timings in milliseconds.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**operation** (STRING): Operation that was completed, such as switchover or switchback.

**subPhaseTimes** (STRING): Report of sub-phase times for the WAFL portion of the MetroCluster operation.

**waf1.mcc.total.maxvvols****Severity**

NOTICE

**Description**

This message occurs in a MetroCluster(tm) configuration when the system cannot configure a FlexVol® volume because the maximum number of volumes has been reached.

**Corrective Action**

Check whether there exists a left-behind aggregate on this node from a previous switchover or takeover operation. If such an aggregate exists, then perform the corrective actions specified in the earlier EMS messages for the left-behind aggregate to return it to its original owner. Otherwise, either destroy some existing volumes or perform a 'volume move' operation to move some existing volumes to a storage aggregate located on a different node to make room for the new ones.

## Syslog Message

Cannot configure FlexVol volume '%s' (UUID %s) in aggregate '%s': There are too many FlexVol volumes on the system.

## Parameters

**vol** (STRING): FlexVol volume name.  
**uuid** (STRING): FlexVol volume unique ID.  
**aggregate** (STRING): Aggregate name.

## wafL.mcc.volLimit.exceeded

### Deprecated

Deprecated as of version 9.2 as the additional volume count soft limit for MCC is removed. The regular MCC limit based on the setting of node-object-limit is still in place.

### Severity

ERROR

### Description

This message occurs when the total volume count exceeds the limit supported on this node.

### Corrective Action

Bring the number of volumes below the supported limit on this node, either by performing a 'volume move' operation to move the extra volumes to a storage aggregate located on a different node, or by destroying some or all of the unused extra volumes on this node.

## Syslog Message

There are more volumes (%d) than the maximum number of volumes (%d) supported on this node for host: %s.

## Parameters

**numVols** (INT): Current number of volumes.  
**limit** (INT): Maximum number of volumes supported.  
**host** (STRING): Host on this node for which the volume limit has been exceeded.

## wafL.mem events

### wafL.mem.stomper

### Severity

NOTICE

### Description

This event indicates whether memory stomper code is running or not.

### Corrective Action

(None).

## Syslog Message

You are %s running in corruption detection diagnostic mode with degraded performance.

## Parameters

**onoff** (STRING): The stomper code is on/off

# wafl.memory.events

## wafl.memory.status

### Severity

INFORMATIONAL

### Description

This message occurs periodically to inform the administrator about the amount of WAFL memory that is currently available.

### Corrective Action

(None).

### Syslog Message

%dMB of memory is currently available for the WAFL file system.

### Parameters

**mem** (INT): Available WAFL memory, in megabytes (MB).

## wafl.memory.statusLowMemory

### Severity

NOTICE

### Description

This message occurs when WAFL is low on memory.

### Corrective Action

(None).

### Syslog Message

WAFL is running low on memory, with %dMB remaining.

### Parameters

**mem** (INT): Available WAFL memory, in megabytes (MB).

## wafl.memory.statusVeryLowMemory

### Severity

ALERT

### Description

This message occurs when WAFL is very low on memory.

### Corrective Action

(None).

## Syslog Message

WAFL is running very low on memory, with %dMB remaining.

## Parameters

**mem** (INT): Available WAFL memory, in megabytes (MB).

# wafl.metadata events

## wafl.metadata.rsrv.decreased

### Severity

NOTICE

### Description

This message occurs when the space reserved for volume metadata is decreased when the value of "space-guarantee" for a FlexClone volume is changed from "volume" to "none". After the decrease in the metadata reserve, the freed up space is added to the space available in the aggregate.

### Corrective Action

(None).

### Syslog Message

Metadata reserve for volume "%s" has decreased.

### Parameters

**volume\_name** (STRING): Volume name.

## wafl.metadata.rsrv.increased

### Severity

NOTICE

### Description

This message occurs when the space reserved for volume metadata is increased when the value of "space-guarantee" for a FlexClone volume is changed from "none" to "volume" or when a FlexClone volume is split from its parent volume. This metadata reserve is required to store volume metadata inside an aggregate.

### Corrective Action

(None).

### Syslog Message

Metadata reserve for volume "%s" has increased.

### Parameters

**volume\_name** (STRING): Volume name.

# wafl.metafile events



## **waf1.metafile.reg.unknown**

### **Severity**

NOTICE

### **Description**

This message occurs when a metafile to be created is not documented, has information about it missing, or has not been identified and associated with an entry from the wiki page: [MetafileRegistration](#).

METAFILE\_UNKNOWN is used as a placeholder in these cases.

### **Corrective Action**

If you are familiar with this metafile, its purpose, or others who have worked with this file, please contact [dl-replication-developer](#).

### **Syslog Message**

%s is declared as METAFILE\_UNKNOWN. If you are familiar with this metafile, its purpose, or others who have worked with this file, please contact [dl-replication-developer](#).

### **Parameters**

**path** (STRING): Filepath of the metafile

## **waf1.metafile.reg.unused**

### **Severity**

NOTICE

### **Description**

This message occurs when a metafile believed to be unused or obsolete is about to be created.

### **Corrective Action**

The information in the wiki page: [MetafileRegistration](#) and [metafiles.h](#) needs to be corrected. Please contact [dl-replication-developer](#).

### **Syslog Message**

%s is declared as METAFILE\_UNUSED. The information in the wiki page: [MetafileRegistration](#) and [metafiles.h](#) needs to be corrected. Please contact [dl-replication-developer](#).

### **Parameters**

**path** (STRING): Filepath of the metafile

## **waf1.mirror.events**

### **waf1.mirror.access.denied**

#### **Severity**

ERROR

#### **Description**

This message occurs when the system rejects client access to a primordial volume. A primordial volume is a volume that has been configured as the destination of a SnapMirror® relationship but has not yet been initialized through that relationship.

## Corrective Action

You can avoid this event by using the 'volume unmount' command to unmount the volume from the namespace, or by using the 'snapmirror initialize' command to initialize the SnapMirror relationship to the volume.

## Syslog Message

Access to SnapMirror volume '%s' was denied because SnapMirror initialization was not complete (exported Snapshot copy identifier '%d' is invalid).

## Parameters

**vol** (STRING): Volume name.  
**snapid** (INT): Exported Snapshot(tm) copy identifier.

## wafl.mirror.upgrade.fail

### Severity

ERROR

### Description

This message occurs when WAFL® unexpectedly encounters a Data ONTAP® 10.0.x mirrored volume during an upgrade. This mirrored volume should have been deleted or broken off from its mirroring relationship prior to the upgrade.

### Corrective Action

Avoid this event by destroying all Data ONTAP 10.0.x mirrored volumes prior to the upgrade. This event resulted in the volume being marked offline, corrupt, and unrecoverable. Because the current version of Data ONTAP does not recognize this mirrored volume, the volume cannot be recovered by file system repair mechanisms such as wafiron and cannot be brought online.

### Syslog Message

Data ONTAP 10.0.x mirror (type %d) cannot be upgraded by the current version of Data ONTAP. Mirrored volume '%s' was marked offline, corrupt, and unrecoverable. Read the upgrade notes for this release. Such mirrored volumes should be deleted or broken off from their mirroring relationships prior to the upgrade.

### Parameters

**type** (INT): Mirror type in Data ONTAP 10.0.x.  
**vol** (STRING): Volume name.

## wafl.mount events

### wafl.mount.inconsistent

#### Severity

ERROR

#### Description

This message occurs when the FlexVol® volume cannot be mounted; the volume is marked as inconsistent.

#### Corrective Action

Unable to mount the volume. Run wafiron to diagnose the volume.

## Syslog Message

WAFL: Unable to mount volume %s due to %s. Volume is marked as inconsistent.

## Parameters

**vvolname** (STRING): Volume name.

**reason** (STRING): Reason for the inability to mount the volume.

## wافل.mount.info

### Severity

NOTICE

### Description

mount message

### Corrective Action

(None).

## Syslog Message

%s

## Parameters

**msg** (STRING): Message from wافل vol mount path which prints how many messages were sent to WAFL.

## wافل.mount.transient.error

### Severity

ERROR

### Description

This message occurs when the FlexVol® volume cannot be mounted due to transient errors; the volume is taken offline.

### Corrective Action

First, attempt to bring the volume online manually by using the "volume online" command. If the issue persists, determine whether the volume is a NetApp® Aggregate Encryption (NAE) volume by using the "volume show -encryption-type aggregate" command. If so, remount the aggregate of the volume by taking it offline (using the "aggr offline" command), and then bringing it online again (using the "aggr online" command). If the issue still persists, Contact NetApp technical support.

## Syslog Message

WAFL: Unable to mount volume %s, UUID %s due to %s. Volume is taken offline due to transient errors.

## Parameters

**vvolname** (STRING): Volume name.

**uuid** (STRING): UUID of the volume.

**reason** (STRING): Reason for the failure to mount the volume.

# wafl.mp events

## wafl.mp.vldb.update.failed

### Severity

ERROR

### Description

This message occurs when the value of the "multipart enabled" setting in the RAID label changes during the restoration of a Snapshot copy, and an attempt to update the value in the Volume Location Database (VLDB) to match the new RAID label value fails.

### Corrective Action

Perform the following steps using the 'debug vreport' command to resolve the inconsistency in the VLDB. 1. Run the "debug vreport show" command to show the object name. 2. Run the "debug vreport fix -type volume -object object-name" command.

### Syslog Message

Updating the VLDB "granular-data" property of volume "%s" with DSID "%08x" to "%s" failed. Reason: %s.

### Parameters

**vol** (STRING): Name of the volume.  
**dsid** (LONGINT): Unique identifier (DSID) of the volume.  
**value** (STRING): New value to be set.  
**errmsg** (STRING): Error message.

# wafl.needs events

## wafl.needs.blog

### Severity

ERROR

### Description

This message occurs when `batched_free_log` needs to be turned on to allow performing one of the requested actions, such as "Dense Undo", "Dense Revert", and "Clone Split".

### Corrective Action

Contact NetApp technical support..

### Syslog Message

`batched_free_log` needs to be turned on.

### Parameters

(None).

# wafl.nfsv4 events

## **waf1.nfsv4.acl.not.mapped**

### **Severity**

ERROR

### **Description**

This message occurs when the system cannot map an ID stored in the Access Control Entry (ACE) to a valid UNIX® User ID/Group ID (UID/GID) pair. This failure makes this ACE an invalid one, resulting in a possible loss of access to the file or directory for which this ACE was set.

### **Corrective Action**

Contact NetApp technical support.

### **Syslog Message**

ID in NFS Access Control Entry (ACE) on inode number %x, file system ID %x could not be mapped to a user or group. Possible loss of access can result.

### **Parameters**

**inum** (INT): Inode number.

**fsid** (INT): File system ID.

## **waf1.no events**

## **waf1.no.large.vvol.size.set**

### **Severity**

ERROR

### **Description**

This message occur when a volume don't have 'large\_size\_enabled' flag set even when the nominal size crosses the 100TB limit.

### **Corrective Action**

(None).

### **Syslog Message**

Nominal size of the "%s" volume exceeded either 100TB limit.

### **Parameters**

**vvolname** (STRING): Name of the volume.

## **waf1.ntfsunixsecurityops events**

## **waf1.ntfsUnixSecurityOps**

### **Severity**

NOTICE

### **Description**

This message occurs when a UNIX-type permissions or ownership change is attempted by an NFS client

on an NTFS (Windows) volume, but changes of this type are prohibited by an export policy rule.

### Corrective Action

Use the "vserver show -fields uuid" command to find the Vserver name corresponding to the Vserver's UUID. Use the "volume file show-inode" command with the file ID and volume name information to find the file or directory path and consider using Windows client to set file or directory attributes.

### Syslog Message

NFS client attempted prohibited change of UNIX permissions for file ID: %d, on NTFS volume : %s%s%s

### Parameters

**fileid** (INT): File ID of the entry in the directory.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Vserver's Universal Unique Identifier (UUID).

## wafn.vlog events

### wafn.vlog.stats.createError

#### Severity

NOTICE

#### Description

This event is issued to indicate that the system is unable to create nvlog stats file.

#### Corrective Action

(None).

#### Syslog Message

Can't create nvlog replay stats file

#### Parameters

(None).

### wafn.vlog.stats.writeError

#### Severity

NOTICE

#### Description

This event is issued to indicate the system is unable to write to nvlog stats file.

#### Corrective Action

(None).

#### Syslog Message

Nvlog replay stats fio\_write failed

## Parameters

(None).

## wafl.nvlog.versionFail

### Severity

ALERT

### Description

This message occurs when the nvlog version found in the NV log is not recognized by the current version of Data ONTAP®.

### Corrective Action

This might be the result of an upgrade or revert attempt after a dirty boot. If so, try the upgrade or revert after a clean shutdown. If the problem persists, or if this was not due to an upgrade or revert, contact NetApp technical support.

### Syslog Message

Could not replay WAFL NV log: nvlog format is of a prior or future release (%d, %d).

### Parameters

**version** (INT): WAFL® NV log version in the NV log.

**myversion** (INT): Kernel WAFL NV log version.

## wafl.nvsave events

### wafl.nvsave.deletion.status

### Severity

INFORMATIONAL

### Description

This message occurs when one or more NVSave files are deleted from the local root volume or the partner root volume while in takeover mode, or from the partner node when the partner node is not in takeover mode.

### Corrective Action

(None).

### Syslog Message

Deleted NVSave files corresponding to %s %s with FSID %d, HA policy %d, CP\_count %d %s on %s node. UUID is %s.

### Parameters

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

**vol\_name** (STRING): Volume name.

**fsid** (INT): File system identifier.

**ha\_policy** (INT): High-availability (HA) policy of the parent aggregate.

**CP\_count** (INT): Volume's WAFL® consistency point (CP) count.

**status** (STRING): Whether deletion failed or succeeded.

**node** (STRING): Whether deletion was attempted on a local or partner node.

**UUID** (STRING): UUID of Volume.

## **wافل.nvsave.overwritingstale**

### **Severity**

NOTICE

### **Description**

This message occurs during replay of the WAFL® nvlog, when the system discovers records associated with an offline volume. These records need to be stored in a system file on the root volume for future replay. However, a previously saved replay file already exists for this volume in the root volume. The system overwrites the stale replay file if the Consistency Point (CP) count stored in the file is less than that in nvlog.

### **Corrective Action**

(None).

### **Syslog Message**

Detected stale system nvlog replay file for a nonexistent volume with FSID '%d' in the '%s' root volume for WAFL nvlog %d. This stale file was overwritten because of '%s'.

### **Parameters**

**fsid** (INT): File system ID (FSID) of the missing volume.

**host** (STRING): Owner of the root volume to which the detected nvlog records are successfully stored.

**log** (INT): WAFL NV log section in which nvlog records for the volume are located.

**error** (STRING): Reason for overwriting the existing nvsave file.

## **wافل.nvsave.replaying**

### **Severity**

NOTICE

### **Description**

This message occurs when a volume comes online and detects relevant nvlog data to be replayed from the system file in the root volume.

### **Corrective Action**

(None).

### **Syslog Message**

Relevant nvlog data detected for volume %s%s%s%s, FSID %d, CP count %d in WAFL NV log %d.  
Replaying nvlog records for the volume from a system file in the root volume.

### **Parameters**

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Consistency Point (CP) count found in the system nvlog file.

**log** (INT): WAFL nvlog section in which nvlog records for the volume are located.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.



## wafl.nvsave.replaying.fail

### Severity

EMERGENCY

### Description

This message occurs when a volume comes online and fails to replay relevant nvlog data from the system file in the root volume. The nvlog data for the volume might be lost.

### Corrective Action

Take appropriate recovery action to recover from data loss, if any.

### Syslog Message

Relevant nvlog data detected for volume %s%s%s%s, FSID %d, CP count %d. Failed to replay nvlog records in WAFL nvlog %d for the volume from a system file in the root volume due to error: %s.

### Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Consistency Point (CP) count found in the system nvlog file that could not be replayed.

**log** (INT): WAFL® nvlog section in which nvlog records for the volume are located.

**error** (STRING): Reason for the failure.

**log0\_replay\_time** (LONGINT): Milliseconds to replay nvlog records from the first NVSave file.

**log0\_cp\_time** (LONGINT): Milliseconds to complete the WAFL CP after nvlog records from the first NVSave file are replayed.

**log1\_replay\_time** (LONGINT): Milliseconds to replay nvlog records from the second NVSave file.

**log1\_cp\_time** (LONGINT): Milliseconds to complete the WAFL CP after nvlog records from the second NVSave file are replayed.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

## wafl.nvsave.replaying.partial

### Severity

NOTICE

### Description

This message occurs when a volume comes online and does a partial replay of the nvlog data from the system file in the root volume. The reason for the partial replay is listed.

### Corrective Action

(None).

### Syslog Message

%s. Nvlog records from the system file in the root volume for volume with FSID %d might have been partially replayed.

### Parameters

**reason** (STRING): Reason for the partial replay.

**fsid** (INT): File system ID (FSID) of the volume.

## waf1.nvsave.replaying.pass

### Severity

NOTICE

### Description

This message occurs when a volume comes online and successfully replays relevant nvlog data from the system file in the root volume.

### Corrective Action

(None).

### Syslog Message

Relevant nvlog data detected for volume %s%s%s%s, FSID %d, CP count %d. Successfully replayed nvlog records for the volume from a system file in the root volume.

### Parameters

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Consistency Point (CP) count found in the system nvlog file.

**log0\_replay\_time** (LONGINT): Milliseconds to replay nvlog records from the first NVSave file.

**log0\_cp\_time** (LONGINT): Milliseconds to complete WAFL® CP after nvlog records from the first NVSave file are replayed.

**log1\_replay\_time** (LONGINT): Milliseconds to replay nvlog records from the second NVSave file.

**log1\_cp\_time** (LONGINT): Milliseconds to complete the WAFL CP after nvlog records from the second NVSave file are replayed.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

## waf1.nvsave.saving

### Severity

NOTICE

### Description

This message occurs during replay of the WAFL® nvlog, when records associated with a nonexistent or currently offline volume are discovered. These records will be stored in a system file on the root volume and will be replayed when the offline volume comes online.

### Corrective Action

(None).

### Syslog Message

Replay detected nvlog records for a nonexistent or offline volume with FSID %d and CP count %d in WAFL nvlog %d. The nvlog records will be stored in a system file in the "%s" root volume for possible future replay.

### Parameters

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Volume's WAFL Consistency Point (CP) count found in the nvlog.

**log** (INT): WAFL nvlog section in which nvlog records for the volume are located.

**host** (STRING): Owner of the root volume on which the detected nvlog records are stored.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

## **wافل.nvsave.saving.fail**

### **Severity**

EMERGENCY

### **Description**

This message occurs when, during replay of the WAFL® nvlog, records associated with a non-existent or offline volume are discovered, but not successfully saved to a system file on the root volume. Data loss is possible when the volume comes back online.

### **Corrective Action**

Nvlog records for an offline volume were not successfully saved, data loss is possible when the volume comes back online. Take appropriate recovery action to recover from data loss, if any, when the volume comes back online.

### **Syslog Message**

Replay detected nvlog records for a nonexistent volume with FSID %d and CP count %d in WAFL nvlog %d. Failed to store nvlog records in a system file in the "%s" root volume due to error: %s.

### **Parameters**

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Volume's WAFL Consistency Point (CP) count found in the nvlog.

**log** (INT): WAFL nvlog section in which nvlog records for the volume are located.

**host** (STRING): Owner of the root volume on which the detected nvlog records were intended to be stored.

**error** (STRING): Error code.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

## **wافل.nvsave.saving.pass**

### **Severity**

NOTICE

### **Description**

This message occurs when, during replay of the WAFL® nvlog, records associated with a currently offline volume are discovered and successfully stored in a system file on the root volume. These records are replayed when the volume comes back online.

### **Corrective Action**

The nvlog records for an offline volume are successfully saved on the root volume. These records are replayed when the volume comes back online.

### **Syslog Message**

Replay detected nvlog records for an offline volume with FSID %d and CP count %d in WAFL nvlog %d. The nvlog records were successfully stored in a system file on the "%s" root volume in %llu milliseconds.

### **Parameters**

**fsid** (INT): File system ID (FSID) of the missing volume.

**cp\_count** (INT): Volume's WAFL Consistency Point (CP) count found in the system nvlog file.

**log** (INT): WAFL NV log section in which nvlog records for the volume are located.

**host** (STRING): Owner of the root volume on which the detected nvlog records are successfully stored.

**timetaken** (LONGINT): Time taken to save the nvlog records, in milliseconds.

**vol\_type** (STRING): Whether this message refers to NVSave files for an aggregate or a flexible volume.

## wafl.objstore events

### wafl.objstore.ownership.conflict

#### Severity

ALERT

#### Description

This message occurs when FabricPool aggregate and SnapMirror® tries to share the same underlying objectstore connection pools in a platform where the sharing is not supported.

#### Corrective Action

For more information or assistance, contact NetApp technical support.

#### Syslog Message

FabricPool aggregate and SnapMirror® cannot share the underlying objectstore connection pool.  
Objectstore Config id: %u, current owner: %u.

#### Parameters

**primary\_config\_id** (INT): Primary object store config id.

**current\_owner** (INT): Existing owner. FabricPool(1)/SnapMirror(2).

## wafl.offline events

### wafl.offline.infoError

#### Severity

ALERT

#### Description

This message occurs when the system attempts to place a volume online, but cannot read the volinfo or fsinfo data for the volume.

#### Corrective Action

Volume is inconsistent. Contact NetApp technical support.

#### Syslog Message

Either both volinfo blocks or fsinfo block on %s %s%s%s%s are/is invalid.

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume, for cases when the volume name is insufficient.

## **wafloffline.openfiles**

### **Severity**

NOTICE

### **Description**

This message occurs when waflliron detects a failure while trying to perform an operation because files are open.

### **Corrective Action**

No corrective action is required for this notification event.

### **Syslog Message**

'%s %s' might not be '%s%s' because of '%u' open file'%s'.

### **Parameters**

**voltype** (STRING): Aggregate.

**volname** (STRING): Volume name.

**operation** (STRING): Operation.

**suffix** (STRING): Suffix.

**uses** (INT): Number of open files.

**plural** (STRING): Plural.

## **wafloffline.proc.file**

### **Severity**

NOTICE

### **Description**

This message occurs when waflliron detects a failure while trying to perform an operation because a process has open files.

### **Corrective Action**

No corrective action is required for this notification event.

### **Syslog Message**

The '%s' process has '%u' file'%s' open.

### **Parameters**

**procname** (STRING): Process name.

**openfiles** (INT): Number of open files.

**plural** (STRING): Plural.

## **wafloffline.proc.file.unres**

### **Severity**

NOTICE

### **Description**

This message occurs when waflliron detects a failure while trying to perform an operation because a process has unresolved open files.

### **Corrective Action**

No corrective action is required for this notification event.

### **Syslog Message**

Could not resolve '%u' of '%s's' open file names.

### **Parameters**

**openfiles** (INT): Number of open files.

**procname** (STRING): Process name.

## **wabl.offline.quiesce.aborted**

### **Severity**

NOTICE

### **Description**

This message occurs when wabliron detects a failure while waiting for a volume reference lock to be taken.

### **Corrective Action**

No corrective action is required for this notification event.

### **Syslog Message**

The wait for uses of '%s %s' to quiesce was aborted.

### **Parameters**

**voltype** (STRING): Aggregate.

**aggrname** (STRING): Aggregate name.

## **wabl.offline.quiesced**

### **Severity**

NOTICE

### **Description**

This message occurs when wabliron detects a success while waiting for a volume reference lock to be taken.

### **Corrective Action**

No corrective action is required for this notification event.

### **Syslog Message**

Uses of '%s %s' preventing '%s' have quiesced.

### **Parameters**

**voltype** (STRING): Aggregate.

**aggrname** (STRING): Aggregate name.

**operation** (STRING): Operation.

## **wabl.offline.versionMismatch**

## Severity

EMERGENCY

## Description

This message occurs when the system is inspecting an offline volume by reading the super block and discovers that there is a version mismatch.

## Corrective Action

Confirm that the correct Data ONTAP® version is in use; if not, reboot or upgrade to the correct version as necessary. For more information or assistance, contact NetApp technical support.

## Syslog Message

This version of Data ONTAP does not recognize the filesystem of offline %s %s%s%s%s. The volume is probably from a later version of the software.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

## wafloffline.volref.msg

## Severity

NOTICE

## Description

This message indicates the volume reference message that prevents the operation on an aggregate.

## Corrective Action

No corrective action is required for this notification event.

## Syslog Message

'%s' prevents '%s' of '%s %s'.

## Parameters

**volrefext** (STRING): Volume reference message text.

**operation** (STRING): Operation.

**voltype** (STRING): Aggregate.

**aggrname** (STRING): Aggregate name.

## wafloodemand events

### wafloodemand.aggr.rsv.fail

## Severity

ERROR

## Description

This message occurs when Data ONTAP® fails to reserve space for a file operation due to lack of space on the aggregate.

## Corrective Action

Create free space on the aggregate by either adding more disks or manually deleting data. If this condition persists, contact NetApp technical support for further assistance.

## Syslog Message

Insufficient space on aggregate '%s %s' to perform file operation on volume '%s%s%s' (failed to re-create %s flexible space pool).

## Parameters

**owner** (STRING): Owner of the aggregate.

**aggr** (STRING): Name of the aggregate.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**fsptype** (STRING): Type of the flexible space pool.

## waf.ondemand.destroy.pending

### Severity

NOTICE

### Description

This message occurs when a file operation cannot be immediately destroyed because the destination file has been captured in one or more snapshots.

### Corrective Action

Delete the snapshots that have captured the destination file.

### Syslog Message

%s

### Parameters

**destroy\_info** (STRING): Status of the file operation.

## waf.ondemand.done

### Severity

NOTICE

### Description

This message occurs when a file operation is completed.

### Corrective Action

(None).

### Syslog Message

%s



## Parameters

**done\_info** (STRING): Status of the completed file operation.

## **waf.ondemand.dst.fsize.exceeded**

### Severity

NOTICE

### Description

This message occurs when the size of a file or LUN being moved or copied exceeds the maximum file or LUN size allowed by the system. The ondemand operation will fail for that particular file or LUN.

### Corrective Action

(None).

### Syslog Message

%s destination job %s-%d. %s %s%s%s%s, FSID 0x%x, The file operation with file ID %d has failed because the file size exceeded the maximum file size allowed on the current system, which is %lu bytes. The size of the file being copied or moved is %lu bytes.

## Parameters

**job\_type** (STRING): Operation type.

**job\_uuid** (STRING): Unique ID (UUID) of the file operation source job.

**file\_index** (INT): Unique identifier for the file in the source job.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume if the volume name itself is insufficient.

**fsid** (INT): File system identifier of the volume.

**fileid** (INT): File identifier.

**maxsize** (LONGINT): Maximum file size allowed on this system.

**filesize** (LONGINT): Size of the file being copied or moved.

## **waf.ondemand.dstPresent**

### Severity

INFORMATIONAL

### Description

This message occurs when a file shows up in the public name space of destination volume during file operation.

### Corrective Action

(None).

### Syslog Message

%s

## Parameters

**file\_info** (STRING): Destination file's UUID and index.

## **waf1.ondemand.errPaused**

### **Severity**

ALERT

### **Description**

This message occurs when a file or LUN operation is paused due to an error condition.

### **Corrective Action**

Based on the information in the event, resolve the problem that is causing the error condition, if possible. The operation is retried automatically on an hourly interval, or you can unpause the operation by using the appropriate command: file copy modify, file move modify, lun copy resume, or lun move resume. If the error condition persists, contact NetApp technical support.

### **Syslog Message**

%s

### **Parameters**

**error\_info** (STRING): Error information for the paused operation.

## **waf1.ondemand.metafileErr**

### **Severity**

ERROR

### **Description**

This message occurs when the special-purpose metafile that Data ONTAP® creates to store information on file operations cannot be read. The volume has been left offline as a consequence.

### **Corrective Action**

Issue the 'volume online' command to attempt to bring the affected volume back online. If this condition persists, contact NetApp technical support for further assistance.

### **Syslog Message**

OnDemand metafile load error, volume '%s%s%s': %s

### **Parameters**

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**error\_info** (STRING): Error information.

## **waf1.ondemand.repaired.metafile**

### **Severity**

ERROR

### **Description**

This message occurs when an inconsistency is discovered in the OnDemand metafile header or entry. The inconsistency is automatically repaired upon detection.

## Corrective Action

(None).

## Syslog Message

%s certification failed: %s, %s, %s, %s, %s, fbn (%llu), record\_index (%llu)

## Parameters

**fh** (STRING): File handle of ondemand metafile which is repaired.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

**condition** (STRING): The initial metadata validation test that failed.

**source** (STRING): The location of the failing metadata validation test.

**fbn** (LONGINT): File block number (fbn) of the repaired block.

**record\_index** (LONGINT): Index of the record entry in the repaired block.

## waf.ondemand.space.rsv.fail

### Severity

ERROR

### Description

This message occurs when Data ONTAP® fails to reserve space for a file operation due to lack of space on the volume.

### Corrective Action

Manually delete data to create free space on the volume, and then modify the job to unpause the data scanner by using the "volume file [copy|move] modify -job-uuid <job\_uuid> -file-index <file\_index> -scanner -paused false" command. If this condition persists, contact NetApp technical support for further assistance.

### Syslog Message

Insufficient space on '%s%s%s' to perform file operation %s-%d (failed to re-create %s flexible space pool).

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**job\_uuid** (STRING): Unique ID (UUID) of the file operation.

**file\_index** (INT): Unique identifier for the file in the job.

**fsptype** (STRING): Type of the flexible space pool.

## waf.ondemand.srcJobRecFail

### Severity

ERROR

### Description

This message occurs when the system fails to repair a corrupted record in the special-purpose metafile that holds information about file operations.

## Corrective Action

The source job recovery will be reattempted the next time the volume comes online. To explicitly reattempt the recovery, take the source volume offline and then bring it back online. If this error condition persists, contact NetApp technical support.

## Syslog Message

Volume %s%s%s: %s

## Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**result** (STRING): Result of the failed recovery.

## waf.ondemand.srcJobRecovery

### Severity

NOTICE

### Description

This message occurs when the system repairs a corrupted record in the special-purpose metafile that holds information about file operations. If the file operation associated with the recovered source job is still in progress, the file operation can be resumed. Otherwise, the original source file can be recovered to lost+found by using the "file source recover" command.

### Corrective Action

(None).

### Syslog Message

Recovered file operation source job %s-%d in volume "%s%s%s".

### Parameters

**job\_uuid** (STRING): Unique ID (UUID) of the file operation source job.

**file\_index** (INT): Unique identifier for the file in the source job.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## waf.ondemand.srcRecovered

### Severity

INFORMATIONAL

### Description

This message occurs when a file operation source recovery is completed.

### Corrective Action

(None).

### Syslog Message

Recovered %s source job %s-%d to {DSID:%d:%s}.

## Parameters

**type** (STRING): Operation type.

**job\_uuid** (STRING): Unique ID (UUID) of the file operation source job.

**file\_index** (INT): Unique identifier for the file in the source job.

**dsid** (INT): Data set identifier (DSID) of the volume where the recovery was performed.

**path** (STRING): Path to the recovered object.

## wافل.ondemand.started

### Severity

NOTICE

### Description

This message occurs when an OnDemand file or LUN operation is started. The "operation\_info" parameter contains information on the type and identity of the started operation.

### Corrective Action

(None).

### Syslog Message

OnDemand operation started: %s

### Parameters

**operation\_info** (STRING): Initial status of the OnDemand operation.

## wافل.ondemand.volBusy

### Severity

ERROR

### Description

This message occurs when a volume operation fails because the volume is the source or destination of one or more file operations. Continuing with the volume operation could potentially disrupt file operations.

### Corrective Action

Make sure that file operations are making progress, and if they are paused unpause them. Retry the volume operation after the completion of the file operations.

### Syslog Message

Volume %s%s%s can not be offlined, until all pending file operations are completed.

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wافل.online events

## **waf1.online.admin.vol**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a cluster admin volume comes online during MetroCluster™ switchover.

### **Corrective Action**

None.

### **Syslog Message**

Cluster admin volume %s DSID (%08x) in aggregate %s is online.

### **Parameters**

**vol** (STRING): Name of the volume.

**dsid** (LONGINT): Unique identifier (DSID) of the volume.

**aggr** (STRING): Name of the containing aggregate.

## **waf1.online.fail.vmalignment**

### **Severity**

EMERGENCY

### **Description**

This message occurs when a virtual volume cannot be brought online because it has the 'vm-align' option set, which is not supported in this release of Data ONTAP®.

### **Corrective Action**

Install a version of Data ONTAP that supports the 'vm-align' option, available from your vendor.

### **Syslog Message**

Volume '%s%s%s%s' is being taken offline because it has the 'vm-align' option set, which is not supported in this version of Data ONTAP.

### **Parameters**

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## **waf1.online.infoError**

### **Severity**

ALERT

### **Description**

This message occurs when there is a problem reading the volinfo or fsinfo data while trying to mount an online volume.

## Corrective Action

(Call support).

## Syslog Message

Neither volinfo block of %s %s%s%s%s is valid OR the fsinfo block is invalid.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

## waf1.online.mode.corrupted

### Severity

ALERT

### Description

This event is issued when there is an inconsistency in the mode bits that indicate whether this volume is in a clustered or nonclustered environment.

### Corrective Action

Run waf1iron on the containing aggregate to fix up the mode bits.

### Syslog Message

The mode bits that indicate whether the %s '%s' volume is in a clustered or nonclustered environment are inconsistent.

### Parameters

**type** (STRING): Type of object (flexible volume).

**vol** (STRING): Name of the flexible volume.

## waf1.online.notCompatibleVer

### Severity

EMERGENCY

### Description

This message occurs when a volume or aggregate cannot be brought online because SnapLock® is not supported on this version of Data ONTAP®.

### Corrective Action

Install a version of Data ONTAP that is supported with SnapLock, available from your vendor.

### Syslog Message

%s %s%s%s%s is being taken offline. SnapLock is not supported on this version of Data ONTAP.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## **wafl.online.toomany**

### **Severity**

ALERT

### **Description**

This message occurs when an attempt to bring a volume online fails because the maximum number of volumes are already online for this node.

### **Corrective Action**

Take a currently mounted volume on this node offline so that you can bring the new volume online.

### **Syslog Message**

%s %s%s cannot be brought online because the maximum online count has been reached.

### **Parameters**

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

## **wafl.online.unknownVersion**

### **Severity**

ALERT

### **Description**

This message occurs when mounting an online volume reveals that the volume version is not recognized by the current version of Data ONTAP®.

### **Corrective Action**

Confirm that the correct Data ONTAP version is in use; if not, reboot or upgrade to the correct version as necessary. For more information or assistance, contact NetApp technical support.

### **Syslog Message**

This version of Data ONTAP does not recognize the filesystem of %s %s%s%s%. The volume is being taken offline.

### **Parameters**

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

## **wafl.online.unsup.deployment**



## Severity

ALERT

## Description

This message occurs when the system determines that the deployment in which this aggregate was created does not match the current (clustered or nonclustered) deployment of Data ONTAP®.

## Corrective Action

Bring the aggregate online on a node that matches the clustered or nonclustered deployment in which it was created.

## Syslog Message

This deployment of Data ONTAP does not support the file system of '%s %s'. It is being taken offline.

## Parameters

**type** (STRING): Type of object (aggregate).

**vol** (STRING): Name of the aggregate.

## wafl.online.vbnMismatch

### Severity

EMERGENCY

### Description

This message occurs when there is an error in the block number data associated with a volume, for example, if WAFL® and RAID have different values for the size of the volume.

### Corrective Action

Run wafliion on the volume. For more information or assistance, contact NetApp technical support.

### Syslog Message

There is a block number error in either the fsinfo block or disk labels on %s %s%s%s%s. The volume is being taken offline.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

## wafl.onoff events

### wafl.onOff.mirType.inv

#### Severity

ERROR

#### Description

This message occurs when WAFL® attempts to bring online or take offline a mirror volume that does not have a valid mirror-type attribute. This situation might be a direct result of an upgrade to a version of Data

ONTAP® that requires the mirror-type attribute or as a result of a corruption.

### Corrective Action

This condition cannot be fixed by file system repair mechanisms such as wafiron. The data is not lost, but it is inaccessible to the administrator. To access the data, contact NetApp technical support. Alternatively, the administrator can destroy the mirror volume by using the 'volume delete -vserver vservice\_name -volume vol\_name' command.

### Syslog Message

The mirror type associated with the mirror volume %s %s%s%s%s is invalid.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name by itself is insufficient.

## wafl.playlist events

### wafl.playlist.bad\_fh

#### Severity

INFORMATIONAL

#### Description

This event indicates bad playlist file handle

#### Corrective Action

(None).

#### Syslog Message

(None).

#### Parameters

**aggr\_name** (STRING): Aggregate name

**vol\_name** (STRING): Volume name

### wafl.playlist.create.failure

#### Severity

INFORMATIONAL

#### Description

This event indicates that playlist creation failed.

#### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**aggr\_name** (STRING): Aggregate name

**vol\_name** (STRING): Volume name

## waf.playlist.no\_playlist

### Severity

INFORMATIONAL

### Description

This event indicates no playlist was found for a volume

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**aggr\_name** (STRING): Aggregate name

**vol\_name** (STRING): Volume name

## waf.pnso events

### waf.pnso.offload.engine.error

### Severity

ALERT

### Description

This message occurs when hardware offload engine requests failed with EIO error.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

%s [%d]

### Parameters

**msg** (STRING): Error message to be displayed.

**err** (INT): Error code received.

## waf.pvbn events

## **waf1.pvbn.check.done**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued to inform the user that the pvbn check has completed. It lists detailed information on what work was performed.

### **Corrective Action**

This message may be disabled by passing -s to waf1 scan pvbn\_check when starting the scanner.

### **Syslog Message**

PVBN check has completed on %s, snapid %d. inos: %d, pvbns: %llu, unks: %llu, errs: %llu.

### **Parameters**

**vol** (STRING): The volume name

**snapid** (INT): ID of the snapshot that has been completely scanned

**inos** (INT): How many inodes were checked

**pvbns** (LONGINT): How many block pointers were checked

**unks** (LONGINT): How many pvbn unknowns were encountered

**errs** (LONGINT): How many pvbn mismatches were encountered

## **waf1.pvbn.check.found.mismatch**

### **Severity**

EMERGENCY

### **Description**

This event is issued to inform the user that the pvbn check has found a pvbn mismatch.

### **Corrective Action**

This message may be disabled by passing -s to waf1 scan pvbn\_check when starting the scanner.

### **Syslog Message**

PVBN check has found a pvbn mismatch on %s, snapid=%d, %s ino=%d, level=%d, fbn=%llu, index=%d, pvbn=%llu, pvbn2=%llu.

### **Parameters**

**vol** (STRING): The volume name

**snapid** (INT): ID of the snapshot being scanned

**ino\_type** (STRING): The inode type

**inum** (INT): Inum being scanned

**level** (INT): Level of the buftree being scanned

**fbn** (LONGINT): Block number containing unknown

**index** (INT): Index in block that is unknown

**pvbn** (LONGINT): PVBN as found in the buftree

**pvbn2** (LONGINT): PVBN as found in the container indirect

## **waf1.pvbn.check.found.unknown**

**Severity**

NOTICE

**Description**

This event is issued to inform the user that the pvbn check has found a pvbn unknown.

**Corrective Action**

This message may be disabled by passing -s to waf scan pvbn\_check when starting the scanner.

**Syslog Message**

PVBN check has found a pvbn unknown on %s, snapid=%d, %s ino=%d, level=%d, fbn=%llu, index=%d.

**Parameters**

**vol** (STRING): The volume name

**snapid** (INT): ID of the snapshot being scanned

**ino\_type** (STRING): The inode type

**inum** (INT): Inum being scanned

**level** (INT): Level of the buftree being scanned

**fbn** (LONGINT): Block number containing unknown

**index** (INT): Index in block that is unknown

**waf.pvbn.check.no.snapshot****Severity**

INFORMATIONAL

**Description**

This event is issued to inform the user that the snapshot being scanned by the pvbn checker no longer exists.

**Corrective Action**

This message may be disabled by passing -s to waf scan pvbn\_check when starting the scanner.

**Syslog Message**

PVBN check on %s aborted, snapshot with snapid %d is gone.

**Parameters**

**vol** (STRING): The volume name

**snapid** (INT): ID of the snapshot that no longer exists

**wafl.qtree events****wafl.qtree.id.mismatch****Severity**

ERROR

**Description**

This message occurs when a qtree ID derived from a WAFL® buffer does not match the given tree ID in a Snapshot(tm) copy.

## Corrective Action

Delete the Snapshot copy containing the qtree for non SnapMirror® DP volumes. For SnapMirror DP volumes, run wafiron with the "-include-mirrors" option set to true to fix the entry. For example, run the "wafiron start -aggr aggr\_name -include-mirrors true" command, where aggr\_name is the aggregate for the given volume.

## Syslog Message

Qtree ID %d is recorded as ID %d in %s %s%s%s%s in Snapshot copy %d.

## Parameters

**tid** (INT): Source tree identifier.

**id** (INT): Qtree root identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**snapld** (INT): Snapshot copy identifier.

## wafl.qtree.invalid.state

### Severity

ERROR

### Description

This message occurs when the qtree file in a Snapshot(tm) copy is in an invalid state.

## Corrective Action

Delete the Snapshot copy containing the qtree for non SnapMirror® DP volumes. For SnapMirror DP volumes, run wafiron with the "-include-mirrors" option set to true to fix the entry. For example, run the "wafiron start -aggr aggr\_name -include-mirrors true" command, where aggr\_name is the aggregate for the given volume.

## Syslog Message

Qtree file in %s %s%s%s%s is in an invalid state %d in Snapshot copy %d.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**reason** (INT): State of the qtree file.

**snapld** (INT): Snapshot copy identifier.

## wafl.qtree.metafile.corrupt

### Severity

ERROR

## Description

This message occurs when the qtree metafile in a Snapshot(tm) copy is inconsistent.

## Corrective Action

Run wafiron to fix the inconsistent volume. Use the "-include-mirrors" option if the volume is a SnapMirror® DP volume. For example, run the "wafiron start -aggr aggr\_name -include-mirrors true" command, where aggr\_name is the aggregate for the given volume.

## Syslog Message

Qtree metafile inconsistent in volume %s%s%s, Snapshot copy %d.

## Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**snapId** (INT): Snapshot copy identifier.

## wafl.qtree.qsmBreak.base

### Severity

ERROR

### Description

This message occurs when the SnapMirror® information of the qtree is being removed as result of WAFL\_check operation.

### Corrective Action

Use the "snapmirror resync" or "snapmirror initialize" to reestablish the SnapMirror relationship.

### Syslog Message

Breaking SnapMirror qtree %d in %s %s%s%s%s: base Snapshot(tm) copy no longer exists. Use the "snapmirror resync" or "snapmirror initialize" to reestablish the SnapMirror relationship.

### Parameters

**qtree** (INT): Qtree identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafl.qtree.qsmBreak.iron

### Severity

ERROR

### Description

This message occurs when the SnapMirror® information of the qtree is being removed as result of a wafiron operation.

## Corrective Action

Use the "snapmirror resync" or "snapmirror initialize" command to reestablish the SnapMirror relationship.

## Syslog Message

Breaking SnapMirror qtree %d in %s %s%s%s%s: wafiron broke all SnapMirror relationships in %s %s%s%s%s. Use the "snapmirror resync" or "snapmirror initialize" command to reestablish the SnapMirror relationship.

## Parameters

**qtree** (INT): Qtree identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident1** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**type2** (STRING): Type of object (volume or aggregate).

**owner2** (STRING): Volume owner.

**vol2** (STRING): Volume name.

**app2** (STRING): Application UUID.

**volident2** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## waf1.qtree.qsmBreak.wc

## Severity

ERROR

## Description

This message occurs when snapmirror information of the qtree is being removed as result of wack

## Corrective Action

Use the "snapmirror resync" or "snapmirror initialize" command to reestablish the SnapMirror relationship.

## Syslog Message

Breaking SnapMirror qtree %d in %s %s%s%s%s: WAFL\_check broke all SnapMirror relationships in %s %s%s%s%s. Use the "snapmirror resync" or "snapmirror initialize" command to reestablish the SnapMirror relationship.

## Parameters

**qtree** (INT): Qtree identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident1** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**type2** (STRING): Type of object (volume or aggregate).

**owner2** (STRING): Volume owner.

**vol2** (STRING): Volume name.

**app2** (STRING): Application UUID.

**volident2** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.



# wafq.quota events

## wafq.quota.control.inconsistent

### Severity

ERROR

### Description

This message occurs when the quota control file is inconsistent, or when it is consistent but the qtree metadata file refers to a vFiler(tm) unit that no longer exists.

### Corrective Action

Reinitialize quotas for the volume.

### Syslog Message

Quota: /vol/%s%s%s or a quota tree on it referred to a nonexistent vFiler unit. Quota: Quotas have been turned off.

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**valident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafq.quota.control.invalid

### Severity

ERROR

### Description

This message occurs when the system discovers that an invalid quota control file exists.

### Corrective Action

Reinitialize quotas for the volume.

### Syslog Message

Quota: Invalid quota control file on %s %s%s%s%s. Quota: Quotas have been turned off.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**valident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafq.quota.db.createFail

### Severity

NOTICE

## Description

This message occurs when the system cannot create the quota database file.

## Corrective Action

(None).

## Syslog Message

Quota: Could not create the quota database on %s %s%s%s%s (volume might be full).

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wabl.quota.db.missing

### Severity

ERROR

### Description

This message occurs when the system is unable to locate a quota usage database.

### Corrective Action

Reinitialize quotas for the volume.

### Syslog Message

WABL: No quota usage database on %s %s%s%s%s.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wabl.quota.db.noMemory

### Severity

ERROR

### Description

This message occurs when the system cannot allocate the quota usage database during initialization.

### Corrective Action

Add more memory or upgrade the hardware.

### Syslog Message

WABL: Could not obtain memory during load of quota usage database. The quota system was not started on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wافل.quota.db.scanFailed

### Severity

ERROR

### Description

This event is issued when the system tries and fails to start the scanner necessary to complete a "quota on" command.

### Corrective Action

Quota initialization has failed due to lack of resources. Re-running a "quota on" after some time should solve the problem.

### Syslog Message

wافل: quota initialization could not be started on %s %s%s%s%s due to lack of resources. Please wait and re-run "quota on".

## Parameters

**type** (STRING): The type of object (volume/aggregate)  
**owner** (STRING): The volume owner  
**vol** (STRING): The volume name  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wافل.quota.group.exceeded

### Severity

NOTICE

### Description

This event is issued when a group quota has been exceeded on a volume. This event will not be repeated for this group for a set time for until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### Corrective Action

Reduce usage by members of this group on this volume or increase the quota and run 'quota resize'.

### Syslog Message

gid %d: disk quota exceeded on volume %s%s%s%s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

## Parameters

**id** (INT): The group identifier  
**owner** (STRING): The volume owner  
**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**mins** (INT): Number of minutes warnings will be suppressed

## **waf1.quota.groupQtree.correct**

### **Severity**

NOTICE

### **Description**

This message occurs when a group quota has been corrected.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**id** (INT): Group identifier.

**tid** (INT): Qtree identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**vwarn** (STRING): vFiler warning suffix.

## **waf1.quota.groupQtree.exceeded**

### **Severity**

NOTICE

### **Description**

This event is issued when a group quota has been exceeded on a tree. This event will not be repeated for this group and tree for a set time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### **Corrective Action**

Reduce usage by members of this group in this tree or increase the quota and run 'quota resize'.

### **Syslog Message**

gid %d tid %d: disk quota exceeded on volume %s%s%s%s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

### **Parameters**

**id** (INT): The group identifier

**tid** (INT): The tree identifier

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.  
**mins** (INT): Number of minutes warnings will be suppressed

## **wافل.quotل.نل.لرلءفلف**

### **Severity**

NOTICE

### **Description**

This message occurs when the system cannot create the quota names file.

### **Corrective Action**

(None).

### **Syslog Message**

quota: Could not create new quota names file on %s %s%s%s%s.

### **Parameters**

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wافل.quotل.نل.لرلءفلف.ءلءلءفلف**

### **Severity**

NOTICE

### **Description**

This message indicates that the system was unable to set the default ACL on an NTFS volume.

### **Corrective Action**

Set the security style to NTFS again or re-create the volume.

### **Syslog Message**

Unable to set default ACL on NTFS %s %s%s%s%s.

### **Parameters**

**type** (STRING): Type of object (volume/aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## **wافل.quotل.ءلءلءفلف.ءلءلءفلف**

### **Severity**

ERROR

## Description

This message occurs when the system could not turn off quotas on the volume.

## Corrective Action

Use the 'quota off' command to turn off quotas on the volume.

## Syslog Message

Quotas could not be turned off on volume %s%s.

## Parameters

**owner** (STRING): Owner of the volume

**vol** (STRING): Name of the volume

## wapl.quota.qtree.correct

### Severity

NOTICE

### Description

This message occurs when a tree quota has been corrected.

### Corrective Action

(None).

### Syslog Message

(None).

### Parameters

**tid** (INT): Qtree identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.

**vwarn** (STRING): vFiler warning suffix.

## wapl.quota.qtree.exceeded

### Severity

NOTICE

### Description

This event is issued when a tree quota has been exceeded on a volume. This event will not be repeated for this tree for a set amount of time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### Corrective Action

Reduce usage in this tree or increase the quota and run 'quota resize'.

## Syslog Message

tid %d: tree quota exceeded on %s %s%s%s%s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

## Parameters

**tid** (INT): The tree identifier

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**mins** (INT): Number of minutes warnings will be suppressed

**object\_uuid** (STRING): Volume instance UUID of the resource object.

## wافل.quota.sec.change

### Severity

INFORMATIONAL

### Description

This message indicates that the security style for a directory has been changed.

### Corrective Action

(None).

## Syslog Message

security style for /vol/%s%s%s/%s changed from %s to %s

## Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**dir** (STRING): Directory name.

**ostyle** (STRING): Former security style.

**nstyle** (STRING): New security style.

## wافل.quota.user.exceeded

### Severity

NOTICE

### Description

This event is issued when a user quota has been exceeded on a volume. This event will not be repeated for this user and volume for a set time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### Corrective Action

Reduce usage by this user on this volume or increase the quota and run 'quota resize'.

## Syslog Message

uid %d: disk quota exceeded on volume %s%s%s%s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

## Parameters

**id** (INT): The user identifier

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**mins** (INT): Number of minutes warnings will be suppressed

## wافل.quota.user.exceeded.win

### Severity

NOTICE

### Description

This event is issued when a user quota has been exceeded on a volume. This event will not be repeated for this user and volume for a set time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### Corrective Action

Reduce usage by this user on this volume or increase the quota and run 'quota resize'.

### Syslog Message

disk quota exceeded on volume %s%s%s%s for Windows user ID: %s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

## Parameters

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**id** (STRING): Windows user identifier

**mins** (INT): Number of minutes warnings will be suppressed

## wافل.quota.userQtrees.correct

### Severity

NOTICE

### Description

This message occurs when a user quota has been corrected.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**id** (INT): User identifier.

**tid** (INT): Qtrees identifier.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.



**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.  
**vwarn** (STRING): vFiler warning suffix.

## **waf1.quota.userQtree.correct.win**

### **Severity**

NOTICE

### **Description**

This message occurs when a user quota has been corrected and the user ID is a SID.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**id** (STRING): Windows identifier value.  
**tid** (INT): Qtree identifier.  
**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier for the volume in cases in which the volume name alone is insufficient.  
**vwarn** (STRING): vFiler warning suffix.

## **waf1.quota.userQtree.exceeded**

### **Severity**

NOTICE

### **Description**

This event is issued when a user quota has been exceeded on a tree. This event will not be repeated for this user and tree for a set time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### **Corrective Action**

Reduce usage by this user in this tree or increase the quota and run 'quota resize'.

### **Syslog Message**

uid %d tid %d: disk quota exceeded on volume %s%s%s%s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

### **Parameters**

**id** (INT): The user identifier  
**tid** (INT): The tree identifier  
**owner** (STRING): The volume owner

**vol** (STRING): The volume name  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.  
**mins** (INT): Number of minutes warnings will be suppressed

## **waf1.quota.userQtree.exceeded.win**

### **Severity**

NOTICE

### **Description**

This event is issued when a user quota has been exceeded on a tree. This event will not be repeated for this user and tree for a set time or until a 'quota resize' is performed. The amount of time is set by the 'quota logmsg' command.

### **Corrective Action**

Reduce usage by this user in this tree or increase the quota and run 'quota resize'.

### **Syslog Message**

tid %d: disk quota exceeded on volume %s%s%s%s for Windows user ID: %s. Additional warnings will be suppressed for approximately %d minutes or until a 'quota resize' is performed.

### **Parameters**

**tid** (INT): The tree identifier  
**owner** (STRING): The volume owner  
**vol** (STRING): The volume name  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.  
**id** (STRING): Windows user identifier  
**mins** (INT): Number of minutes warnings will be suppressed

## **waf1.raid events**

### **waf1.raid.cft.aggr.snapshot.missing**

#### **Severity**

ERROR

#### **Description**

This message occurs during an aggregate transition workflow when an aggregate Snapshot® copy created in a 7-Mode environment is missing.

#### **Corrective Action**

Return the disks to the 7-Mode environment, create the Snapshot copy for the aggregate, and then repeat the transition workflow.

#### **Syslog Message**

Aggregate transition failed for aggregate %s due to unavailability of the 7-Mode aggregate snapshot.

#### **Parameters**

**aggrname** (STRING): Name of the aggregate.

## waf1.raid.incons.buf

### Severity

ERROR

### Description

This message occurs when the WAFL® kernel detects an inconsistent block on a disk. Access to the block is disabled until corrective action is taken.

### Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. If the volume is inconsistent, run wafiron as soon as possible to clear the inconsistency. For a SnapMirror® target, rebuild from the source or resynchronize. If the block is in a Snapshot copy, delete the Snapshot copy. If the problem persists, contact NetApp technical support for more options.

### Syslog Message

WAFL inconsistent: bad block at VBN %llu (vvbn:%llu fbn:%llu level:%u) in %s inode (fileid:%s snapid:%u fixable:%d file\_type:%d disk\_flags:0x%x error:%u raid\_set:%d) in volume %s%s%s%s.

### Parameters

**vbn** (LONGINT): Volume block number.

**vvbn** (LONGINT): Virtual block number.

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**file\_space** (STRING): File space.

**fileid** (STRING): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**fixable** (INT): Boolean value indicating whether iron can repair this block.

**type** (INT): File type.

**flags** (LONGINT): Inode disk flags.

**error** (INT): Error on the block.

**raid\_set** (INT): Boolean indicating whether RAID set the error.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf1.raid.incons.clear.error

### Severity

NOTICE

### Description

This message occurs when Data ONTAP® resets the error on a data block.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Resetting error (%u) on data block in %s inode (fileid:%u snapid:%u fbn:%llu level:%u) in volume %s%s%s%s

## Parameters

**error** (INT): Error number that was reset.

**inotype** (STRING): Inode space (public or private).

**fileid** (INT): Identifying number for the inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

## wafl.raid.incons.dir.block

### Severity

ERROR

### Description

This message occurs when the WAFL® kernel detects an inconsistent directory block on disk, causing some entries on the directory not to appear. The message indicates the fileid and Snapshot(tm) copy of the block being accessed.

### Corrective Action

The volume has a bad directory block. Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Recover the directory from a backup. For a SnapMirror® target, rebuild from the source or resynchronize. If the block is in a Snapshot copy, delete the Snapshot copy. Then contact NetApp technical support.

### Syslog Message

WAFL inconsistent: bad directory block in inode (fileid:%u snapid:%u) in volume %s%s%s%s.

## Parameters

**fileid** (INT): Identifying number for the suspect inode

**snapid** (INT): Snapshot copy identifier. (0) - active file system. (non-zero) - refers to Snapshot copy.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## wafl.raid.incons.set.err.obj

### Severity

ALERT

### Description

This message occurs when Data ONTAP® marks a block in a capacity tier object with an error. This is useful for tracking movement of pseudobad and unverified blocks.

### Corrective Action

Contact Contact NetApp technical support..

## Syslog Message

Data block at VBN %llu tiered to capacity tier (objid %llu, slot %u, btuuid %s, seqno %llu) and marked with error (%u). Block previously had error (%u).

## Parameters

**vbn** (LONGINT): Volume block number.  
**object\_id** (LONGINT): Object identifier.  
**slot** (INT): Slot number.  
**btuuid** (STRING): Buftree UUID.  
**seqno** (LONGINT): Sequence number.  
**error** (INT): Error on the block.  
**prev\_error** (INT): Previous error on the block.

## waf.raid.incons.set.error

### Severity

ALERT

### Description

This message occurs when Data ONTAP® marks an on-disk data block with an error. This is useful for tracking pseudo bad and unverified blocks.

### Corrective Action

Contact NetApp technical support.

## Syslog Message

Data block in %s inode (fileid: %u snapid: %u fbn: %llu level: %u) in volume %s%s%s%s marked with error (%u) by %s. Block previously had error (%u).

## Parameters

**inotype** (STRING): Inode space (public or private).  
**fileid** (INT): Identifying number for the inode.  
**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy  
**fbn** (LONGINT): File block number.  
**level** (INT): Buffer level.  
**owner** (STRING): Owner of the volume.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.  
**error** (INT): Error on the block.  
**module** (STRING): Module that is marking the block with the error.  
**prev\_error** (INT): Previous error on the block.

## waf.raid.incons.snapmap

### Severity

NOTICE

### Description

This message occurs when WAFL detects a RAID inconsistency error in a snapmap block. The message recommends to delete the snapshot in which inconsistent snapmap resides.

## Corrective Action

(None).

## Syslog Message

Snapshot(snapid %u) block accounting metadata is corrupted. To avoid trapping unnecessary blocks in the snapshot, please delete the snapshot.

## Parameters

**snapid** (INT): Snapshot copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

## wافل.raid.incons.snapmap.chk

## Severity

NOTICE

## Description

This message occurs when the WAFL® kernel detects a RAID inconsistency error in a Snapshot(tm) copy active map block, even though the block content passes a metadata check.

## Corrective Action

Contact NetApp technical support.

## Syslog Message

(None).

## Parameters

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**snapid** (INT): Snapshot copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

## wافل.raid.incons.snapshot

## Severity

ERROR

## Description

This message occurs when WAFL® detects a RAID inconsistency error in a snapshot. The message indicates what snapshots have the inconsistent block, and how to recover from the inconsistency.

## Corrective Action

Delete the specified snapshot(s) during the next reboot, either by using the boot menu command 'boot\_snap\_delete' or by using the CLI command 'vol online -r'.

## Syslog Message

VBN %llu has RAID inconsistency errors on volume %s. This block is in use by one or more snapshots. These snapshots should be deleted in the next boot. The affected snapshot IDs are: %s.

## Parameters

**vbn** (LONGINT): Virtual block number.

**volname** (STRING): Name of the volume.

**snapstring** (STRING): The string containing the list of snapid(s) of the snapshots that have the inconsistent blocks.

## wafl.raid.incons.ssd.cached.buf

### Severity

NOTICE

### Description

This message occurs when the WAFL® kernel detects an inconsistent cached block on a ssd disk in a hybrid aggregate. Any access to this block will result in the io getting redirected to the source location on hdd disk.

### Corrective Action

An error on a cached ssd block will not result in bad data being served to the reader: the io will be redirected to hdd and the correct data will be served. So, no corrective action is necessary for this case. But, if there are many errors on data cached on ssd disks, then the ssd disks may be verified for disk consistency.

### Syslog Message

WAFL inconsistent: bad cached ssd block %llu (vvbn:%llu fbn:%llu level:%u) in %s inode (fileid:%s snapid:%u file\_type:%d disk\_flags:0x%x error:%u) in volume %s%s%s%s. IO will be redirected to the source hdd block %llu.

## Parameters

**vbn** (LONGINT): Volume SSD block number.

**vvbn** (LONGINT): Virtual block number.

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**file\_space** (STRING): File space: private or public.

**fileid** (STRING): Identifying number for the inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**type** (INT): File type.

**flags** (LONGINT): Inode disk flags.

**error** (INT): Error on the block.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**valident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**hdd\_vbn** (LONGINT): Volume HDD block number.

## wafl.raid.incons.userdata

### Severity

ERROR

### Description

This message occurs when the WAFL® kernel detects an inconsistent user data block on a disk. Access to this block is disabled until you recover it from a backup.

## Corrective Action

The volume has an inconsistent user data block. Use the "volume file show-inode" command with the file ID and volume name information to find the file path. Recover the file from a backup. For a SnapMirror® target, rebuild from the source or resynchronize. If the block is in a Snapshot(tm) copy, delete the Snapshot copy. Contact NetApp technical support for more options.

## Syslog Message

WAFL inconsistent: inconsistent user data block at VBN %llu (vvbn:%llu fbn:%llu level:%u) in %s inode (fileid:%s snapid:%u file\_type:%d disk\_flags:0x%x error:%u raid\_set:%d) in volume %s%s%s%s.

## Parameters

**vbn** (LONGINT): Volume block number.

**vvbn** (LONGINT): Virtual block number.

**fbn** (LONGINT): File block number.

**level** (INT): Buffer level.

**file\_space** (STRING): File space.

**fileid** (STRING): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**type** (INT): File type.

**flags** (LONGINT): Inode disk flags.

**error** (INT): Error on the block.

**raid\_set** (INT): Boolean indicating whether RAID set the error.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf.raid.incons.xidata

### Severity

EMERGENCY

### Description

This message occurs when the WAFL® kernel detects an inconsistent extended inode data block on disk.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

WAFL inconsistent: bad extended inode data block at VBN %llu (vvbn:%llu fbn:%llu) in inode (fileid:%u snapid:%u disk\_flags:0x%x error:%u raid\_set:%d) in volume %s%s%s%s.

### Parameters

**vbn** (LONGINT): Volume block number.

**vvbn** (LONGINT): Virtual block number.

**fbn** (LONGINT): File block number.

**fileid** (INT): Identifying number for the suspect inode.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**flags** (LONGINT): Inode disk flags.

**error** (INT): Error on the block.

**raid\_set** (INT): Boolean indicating whether RAID set the error.

**owner** (STRING): Owner of the volume.



**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**inotype** (STRING): Inode space (public or private).

## wafl.rclm events

### wafl.rclm.est.scan.done

#### Severity

NOTICE

#### Description

This message occurs when an attempt to estimate the space that can be reclaimed for the aggregate was completed successfully.

#### Corrective Action

(None).

#### Syslog Message

Estimate Reclaim Scanner done for %s with estimated free logical blocks=%llu, physical blocks=%llu, duration=%ld ms.

#### Parameters

**aggr** (STRING): Aggregate name.

**logical\_free\_blks** (LONGINT): Estimated number of reclaimable logical blocks found by the scanner.

**physical\_free\_blks** (LONGINT): Estimated number of reclaimable physical blocks found by the scanner.

**run\_time** (LONGINT): Duration of the scan in milliseconds.

## wafl.readdir events

### wafl.readdir.expired

#### Severity

ERROR

#### Description

This message occurs when a READDIR file operation has exceeded the timeout that it is allowed to run in WAFL. This can be the result of very large or sparse directories, and corrective action is recommended.

#### Corrective Action

You can find information specific to recent directories that have had READDIR file operations expire by using the 'diag' privilege nodeshell CLI command: "wafl readdir notice show" If a directory is indicated as sparse, it is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file. If a directory is not indicated as sparse and the directory is large, it is recommended that you reduce the size of the directory file by reducing the number of file entries in the directory.

#### Syslog Message

A READDIR file operation has expired for the directory associated with volume "%s%s/%s" Snapshot copy ID %d and inode %d.

## Parameters

**volume** (STRING): Volume name that the directory is in.

**app** (STRING): Application universal unique identifier (UUID).

**volident** (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.

**snapid** (LONGINT): Snapshot(tm) copy ID that the directory is in.

**directory\_inum** (LONGINT): The inode number (file ID) of the directory.

**object\_uuid** (STRING): Volume instance UUID of the resource object.

## waf1.readdir.sparse

### Severity

NOTICE

### Description

This message occurs when a REaddir file operation detects that a directory is sparse while iterating over its directory entries. Iterating over sparse directories could negatively impact performance on the node.

### Corrective Action

You can find information specific to directories that have had REaddir file operations detect a sparse directory by using the 'diag' privilege nodeshell CLI command: "waf1 readdir notice show" It is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file.

### Syslog Message

A REaddir file operation has detected that a directory associated with volume %s%s/%s Snapshot ID %d and inode %d is sparse.

## Parameters

**volume** (STRING): Volume name that the directory is in.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.

**snapid** (LONGINT): Snapshot(tm) copy ID the that the directory is in.

**directory\_inum** (LONGINT): The inode number (fileid) of the directory.

## waf1.reallocate events

### waf1.reallocate.aggr.interrupt

### Severity

INFORMATIONAL

### Description

This message occurs when an aggregate reallocation scan is interrupted. The scan was scheduled using the 'reallocate -A' command.

### Corrective Action

None required.

## Syslog Message

Aggregate reallocation scan on '%s' was interrupted.

## Parameters

**path** (STRING): Scan path.

## waf1.reallocate.check.high

### Severity

INFORMATIONAL

### Description

This event is issued if an allocation check scan (issued before reallocation starts) returns indicating that the path requires reallocation.

### Corrective Action

None required. The reallocation scan will automatically be started.

## Syslog Message

Allocation check on '%s' is %d, hotspot %d (threshold %d), so will reallocate.

## Parameters

**path** (STRING): The scan path

**optim** (INT): The current optimization of the path

**hot\_spot** (INT): The current hot-spot optimization of the path

**threshold** (INT): The reallocation threshold of the path

## waf1.reallocate.check.highAdvise

### Severity

INFORMATIONAL

### Description

This event is issued if an allocation check scan returns indicating that the path would benefit from reallocation.

### Corrective Action

Consider reallocating the path (see na\_reallocate(1))

## Syslog Message

Allocation check on '%s' is %d, hotspot %d (threshold %d), consider running reallocate.

## Parameters

**path** (STRING): The scan path

**optim** (INT): The current optimization of the path

**hot\_spot** (INT): The current hot-spot optimization of the path

**threshold** (INT): The reallocation threshold of the path

## waf1.reallocate.check.value

**Severity**

INFORMATIONAL

**Description**

This event is issued to report the optimization of a path.

**Corrective Action**

None required.

**Syslog Message**

Allocation measurement check on '%s' is %d.

**Parameters**

**path** (STRING): The scan path

**optim** (INT): The path optimization.

**wافل.allocate.full.done****Severity**

INFORMATIONAL

**Description**

This event is issued to report the completion of a full, one-time, reallocation scan.

**Corrective Action**

None required.

**Syslog Message**

Full reallocation on '%s' is complete.

**Parameters**

**path** (STRING): The scan path

**wافل.allocate.full.interrupt****Severity**

INFORMATIONAL

**Description**

This event is issued to report that a full, one-time, reallocation scan was interrupted.

**Corrective Action**

None required.

**Syslog Message**

Full reallocation on '%s' was interrupted.

**Parameters**

**path** (STRING): The scan path

## **waf1.reallocate.path.invalid**

### **Severity**

ERROR

### **Description**

This event is issued if an invalid scan path is detected.

### **Corrective Action**

Check why the scan path is invalid, including file removal or volume offline. If the condition is persistent remove the associated reallocation scan.

### **Syslog Message**

Unavailable/invalid reallocation path '%s' detected.

### **Parameters**

**path** (STRING): The scan path

## **waf1.reallocate.recheck.ok**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued if an allocation recheck scan (issued after reallocation has completed) returns indicating the optimization is under the threshold.

### **Corrective Action**

None required.

### **Syslog Message**

Allocation check on '%s' is under threshold.

### **Parameters**

**path** (STRING): The scan path

## **waf1.reallocate.scan.exists**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued if an attempt is made to add a new reallocation scan but a scan job already exists for the supplied path.

### **Corrective Action**

Check the path is correct. If it is correct then delete the previous reallocation job, and add a new one. Use the 'reallocate status' command to list reallocation jobs.

### **Syslog Message**

A reallocation scan for %s already exists.

## Parameters

**path** (STRING): The path with the scan.

## wافل.reallocate.scan.noStart

### Severity

NOTICE

### Description

Reallocation tasks, such as measuring the layout of a file or performing block reallocation, may be performed using an asynchronous WAFL scanner. This event is issued if an error occurs unexpectedly while trying to start such a WAFL scanner.

### Corrective Action

(Call support).

### Syslog Message

reallocate: unable to start '%s' scanner for '%s': %d (%s)

### Parameters

**scan\_type** (STRING): The type of scan operation.

**path** (STRING): The scan path or aggregate name.

**err\_code** (INT): A numerical error code indicating the reason why the scanner could not be started.

**err\_string** (STRING): A string interpreting the err\_code.

## wافل.reallocate.scan.remove

### Severity

INFORMATIONAL

### Description

This event is issued when a reallocation job is persistently removed from the system.

### Corrective Action

None required.

### Syslog Message

Reallocation scan for '%s' is being removed.

### Parameters

**path** (STRING): The path with the scan.

## wافل.reallocate.schedule.invalid

### Severity

ERROR

### Description

This event is issued if an invalid reallocation schedule is detected.

### Corrective Action

Correct the schedule syntax then retry the operation

### Syslog Message

Invalid schedule: %s

### Parameters

**reason** (STRING): The reason why the schedule is invalid

## wafl.relocate events

### wafl.relocate.statusvol.fail

#### Severity

ERROR

#### Description

This message occurs when the system cannot perform an aggregate relocation because one or more of its volumes are used by wafiron to host status files. The relocation is aborted.

#### Corrective Action

Use "aggr wafiron show" command to monitor the status of wafiron and retry aggregate relocation after wafiron is complete.

#### Syslog Message

%s %s%s%s%s is being used by wafiron to host status files. Canceling aggregate relocation.

#### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

### wafl.relocate.vetoedByNvsave

#### Severity

ERROR

#### Description

This message occurs when an aggregate relocation is vetoed because there are NVSave files for an aggregate or a flexible volume that must be replayed before the aggregate or flexible volume can be relocated.

#### Corrective Action

Retry the aggregate relocation after the NVSave files either are replayed (by bringing the aggregate or flexible volume online) or deleted from the root volume. Successful replay, for example, is indicated by the EMS message 'wafl.nvsave.replaying.pass'.

## Syslog Message

Aggregate relocation vetoed because NVSave files are present for %s FSID %d name %s%s%s. UUID is %s.

## Parameters

**vol\_type** (STRING): Whether this message applies to an aggregate or a flexible volume.

**fsid** (INT): File system identifier.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**valident** (STRING): Unique identifier of the volume in cases in which the volume name itself is insufficient.

**UUID** (STRING): UUID of the volume.

# wafl.relocating events

## wafl.relocating.fence.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform an aggregate relocation because a fence is set on one of the volumes in the aggregate. The relocation is aborted.

### Corrective Action

A volume in the aggregate is currently fenced for group operations. Retry aggregate relocation after any group SnapRestore® or mirror transfers are complete. If the problem persists, contact technical support.

## Syslog Message

%s %s%s%s%s has a fence set. Canceling aggregate relocation.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**valident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.relocating.iron.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform an aggregate relocation because wafiron is active. The relocation is aborted.

### Corrective Action

Retry aggregate relocation after wafiron is finished.



## Syslog Message

%s %s%s%s%s is running vol wafiron. Canceling aggregate relocation.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.relocating.sp.fail

### Severity

ERROR

### Description

This message occurs when the system cannot perform an aggregate relocation because a secure purge operation is in progress on one of the volumes in the aggregate. The relocation is aborted.

### Corrective Action

A volume in the aggregate is currently going through a secure purge operation. Retry aggregate relocation after this operation is complete, or after you abort the operation by using the "volume encryption secure-purge abort" command. If the problem persists, contact NetApp technical support.

## Syslog Message

A secure purge operation is in progress on %s %s%s%s. Canceling aggregate relocation.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## wafl.repl events

### wafl.repl.badTimestamp

### Severity

NOTICE

### Description

This message occurs when a replication operation receives a timestamp with an invalid nanoseconds portion, and clears it.

### Corrective Action

(None).

## Syslog Message

Encountered invalid %s time on inode %u of volume %u.

## Parameters

**type** (STRING): Type of timestamp.  
**inode\_num** (INT): Identifier of the affected file.  
**fsid** (INT): Identifier of the volume.

## wافل.repl.maxdirsize.notice

### Severity

NOTICE

### Description

This message occurs when WAFL® sets the maximum directory size of a SnapMirror® destination volume to a value higher than its default during a SnapMirror transfer.

### Corrective Action

(None).

### Syslog Message

The maxdirsize (%u KB) of volume %s (FSID %u) is higher than the default (%u KB). There might be a performance penalty when doing operations on large directories. This was generated for a SnapMirror relationship with ID %s.

## Parameters

**current\_maxdirsize** (INT): Current maximum directory size value.  
**vol** (STRING): Name of the volume.  
**volFsid** (INT): File system ID (FSID) of the volume.  
**default\_maxdirsize** (INT): Default maximum directory size value.  
**SnapMirrorID** (STRING): Unique ID of the SnapMirror relationship for this operation, which can be used to match this message with messages for the same relationship in other logs (for example, the SnapMirror log at /mroot/etc/log/snapmirror\_audit).

## wافل.repl.src.inconsistent

### Severity

ERROR

### Description

This message occurs when the replication source engine encounters a disk error in the file system metadata. Replication transfer will abort.

### Corrective Action

Run wافلiron to fix the active file system. To continue replication transfer, delete the Snapshot(tm) copies containing this bad block.

### Syslog Message

%s: %s%s%s%s: WAFL metadata inconsistent due to disk errors. (vvbn:%u level:%u file\_type:%d).  
Aborting transfer.

## Parameters

**vol\_type** (STRING): Whether the block was being accessed for a volume or an aggregate.  
**owner** (STRING): String indicating the owner of the affected volume or aggregate.  
**vol** (STRING): Name of the volume or aggregate.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**vvbn** (INT): Virtual Volume Block Number.

**level** (INT): Buffer level.

**file\_type** (INT): File type.

## **waf1.repl.zombie.scan.susp**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a logical break zombie restart is going through a scan of composite zombies and it suspends the incoming client composite delete requests.

### **Corrective Action**

(None).

### **Syslog Message**

The deletion of file with fileid %d is suspended because volume %s (FSID %u) is going through a logical break. The file delete will be resumed after logical break is complete.

### **Parameters**

**fileid** (INT): File identifier.

**vol** (STRING): Name of the volume.

**volFsid** (INT): File system ID (FSID) of the volume.

## **waf1.replay.events**

### **waf1.replay.aggr.early.mt**

#### **Severity**

NOTICE

#### **Description**

This message occurs when a partner aggregate is mounted early before replay.

#### **Corrective Action**

(None).

#### **Syslog Message**

Partner aggregate %s is mounted early before WAFL NV log replay.

#### **Parameters**

**aggrname** (STRING): Name of the aggregate.

### **waf1.replay.completed**

#### **Severity**

NOTICE

## Description

This message occurs when WAFL® replay is complete following an abrupt failure, such as a panic. This event logs the time spent in the various phases of WAFL replay.

## Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**log** (STRING): Entry count, bytes used, and duration of operation replay.

**catalog** (STRING): Catalog statistics for WAFL replay.

**swarm** (STRING): Information about the cost of operations during WAFL replay.

**nvce** (STRING): Information about number of nvce hits, skips and misses during WAFL replay.

## wافل.replay.early.mounted

### Severity

NOTICE

### Description

This message occurs when partner aggregates are mounted early before replay.

### Corrective Action

(None).

### Syslog Message

%d of %d partner aggregates are mounted early before WAFL NV log replay.

### Parameters

**early\_aggrs** (INT): Number of aggregates mounted early.

**total\_aggrs** (INT): Total number of data aggregates being replayed.

## wافل.replay.nvlog.ignore.entries

### Severity

NOTICE

### Description

This message occurs when entries are being ignored during replay.

### Corrective Action

(None).

### Syslog Message

%s Ignoring entries in %s log %d for %s %s%s%s%s FSID 0x%x. Volume CP count:%d, Log count: %d.

### Parameters

**reason** (STRING): Reason for ignoring entries.

**log\_host** (STRING): Host in replay.

**log** (INT): Log section identifier.  
**voltype** (STRING): Volume type: volume or aggregate.  
**volowner** (STRING): Owner of the volume: local or partner.  
**app** (STRING): Application UUID.  
**volname** (STRING): Name of the volume or aggregate.  
**server\_uuid** (STRING): Vserver's Universal Unique Identifier (UUID), if any.  
**FSID** (INT): File system identifier.  
**vol\_cp\_count** (INT): Consistency Point (CP) count on the volume.  
**log\_cp\_count** (INT): CP count found in the nvlog.

## **waf1.replay.nvlog.not.needed**

### **Severity**

NOTICE

### **Description**

This message occurs when all entries in a log were ignored, so no replay was performed.

### **Corrective Action**

(None).

### **Syslog Message**

No replay was required for %s log %d.

### **Parameters**

**log\_host** (STRING): Host in replay.  
**log** (INT): Log section identifier.

## **waf1.replay.resource.limited**

### **Severity**

NOTICE

### **Description**

This message occurs when the WAFL replay admin cannot obtain all of the system resources required to allocate the optimal number of replay messages.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**log\_no** (INT): The NVLOG log number.  
**host** (INT): The host in replay.  
**msgs\_wanted** (INT): The number of replay messages requested.  
**msgs\_allocated** (INT): The number of replay messages allocated.

## wafl.replay.swizzle.fsid

### Severity

INFORMATIONAL

### Description

The volume was using the old FSID at the time the operations were recorded in the NVLog. Now that the filer is about to replay logged operations, the filer will use the new FSID for the volume. No corrective action is required.

### Corrective Action

(None).

### Syslog Message

FSID collision mapping created during replay for %s %s%s%s old fsid:%x, new fsid:%x, uuid:%s.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique Identifier of the volume in cases where the volume name itself is insufficient.

**old\_fsid** (INTHEX): Old filesystem identifier

**new\_fsid** (INTHEX): New filesystem identifier

**uuid** (STRING): UUID of the volume

## wafl.root events

### wafl.root.autocreate

#### Severity

NOTICE

#### Description

This event is issued when no online root volume is found. It is used to indicate that auto creation of a root volume will happen.

#### Corrective Action

(None).

#### Syslog Message

No online root volume found on root aggregate %s. Automatically creating one.

#### Parameters

**aggr** (STRING): The root aggregate name

### wafl.root.content.changed

#### Severity

ERROR

## Description

This message occurs when the contents of the root volume are changed to those of a previous point in time. One of the following operations or events can cause the contents of the root volume to change: - The root volume is changed from the boot menu. - The nvlog is jettisoned. - The system is booted with prev\_cp. - The 'wafliroot -prev\_cp' command is used and the changes are committed. - A SnapRestore® operation is done on the root volume.

## Corrective Action

Verify that all recent configuration changes are still in effect.

## Syslog Message

Contents of the root volume '%s' might have changed. Verify that all recent configuration changes are still in effect.

## Parameters

**vol** (STRING): Name of the root volume.

## wafliroot.create

### Severity

INFORMATIONAL

### Description

This message occurs when the system creates a root volume.

### Corrective Action

(None).

### Syslog Message

New root volume %s created.

### Parameters

**vol** (STRING): Name of the root volume.

## wafliroot.mark.error

### Severity

NOTICE

### Description

This event is issued when there is an error in syncing the vvol raid label of a newly created root volume.

### Corrective Action

(None).

### Syslog Message

Error setting volume %s as root. Changes may not be saved.

### Parameters

**vol** (STRING): The root volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wافل.root.overwritesUnsafe**

### **Severity**

ERROR

### **Description**

This message occurs when the root volume is not volume guaranteed or does not guarantee overwrites for space reserved files.

### **Corrective Action**

Set the Space Guarantee on the root volume to "volume" and the Fractional Reserve on the root volume to 100%.

### **Syslog Message**

Root volume %s does not protect overwrites in space reserved files.

### **Parameters**

**vol** (STRING): Root volume name.

## **wافل.s2c events**

### **wافل.s2c.snap.check.failed**

#### **Severity**

ERROR

#### **Description**

This message occurs when SnapMirror® logical transfer is not possible due to incompatible Snapshot copy.

#### **Corrective Action**

Create new Snapshot copy on the source volume and use that for Logical transfer.

#### **Syslog Message**

SnapMirror operation failed in volume '%s%s%s' for Snapshot ID %d because %s.

#### **Parameters**

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**snapId** (INT): Identifier for the Snapshot(tm) copy.

**reason** (STRING): Specific reason for failure.

### **wافل.s2c.snap.lang.mismatch**

#### **Severity**

ERROR



## Description

This message occurs when SnapMirror® restore cannot proceed due to language mismatch at source and destination end.

## Corrective Action

Create new volume with correct language settings.

## Syslog Message

SnapMirror operation failed in %s because of language mismatch between source and destination. Language at destination: '%s'. Language at source: '%s'.

## Parameters

**vol** (STRING): Name of the volume.

**src\_lang\_setting** (STRING): Language setting at source.

**dest\_lang\_setting** (STRING): Language setting at destination.

# wafl.scan events

## wafl.scan.abort.nospace

### Severity

ERROR

### Description

This message occurs when the system is aborting the WORM clear scan because there is no space available in the volume.

### Corrective Action

Create space by increasing the volume or aggregate sizes, or by deleting data or deleting Snapshot copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot copies, use the "volume snapshot delete" command. To increase an aggregate's size, add disks by using the "storage aggregate add-disks" command. Aggregate Snapshot copies are deleted automatically when the aggregate is full.

### Syslog Message

Aborting WAFL %s scanner on volume %s because there is no space available in the volume.

### Parameters

**type** (STRING): String to indicate the scanner that is aborted.

**vol** (STRING): Volume name.

## wafl.scan.autocommit.disable

### Severity

INFORMATIONAL

### Description

This event is issued to inform the user that the system is aborting the autocommit scanner on a Snaplock volume.

## Corrective Action

(None).

## Syslog Message

Aborting autocommit scan on %s %s%s%s%s.

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner of the snaplock volume where the autocommit scanner was aborted.

**vol** (STRING): The name of the snaplock volume where the autocommit scanner was aborted.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the snaplock volume where the autocommit scanner was aborted in cases where volume name itself is insufficient.

## wafI.scan.blkreclaim.activeDelta

### Severity

ERROR

### Description

This event indicates a flexible mirror volume is unable to relinquish free blocks to its containing aggregate. As a result, the containing aggregate may run out of space if this state persists beyond the following incremental transfer.

### Corrective Action

Monitor the state after the next incremental update for this volume and the space available in its containing aggregate

### Syslog Message

Flexible volume '%s%s' cannot release free blocks to its containing aggregate '%s'.

### Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**aggr** (STRING): The aggregate name

## wafI.scan.blockcheck.done

### Severity

NOTICE

### Description

This event is issued when the 'aggr block\_check' tool finishes scanning an aggregate.

### Corrective Action

(None).

### Syslog Message

Writing WAFL aggr block\_check log %s

## Parameters

**filename** (STRING): The name of the log file.

## **wafldb.scan.br.realloc.done**

### Severity

INFORMATIONAL

### Description

This message occurs when a block reallocation scan is complete on an aggregate. The block reallocation scan improves WAFL® write efficiency by reducing free space fragmentation. It runs automatically when aggregate reallocation is scheduled using the 'reallocate -A' command.

### Corrective Action

(None).

### Syslog Message

Block reallocation scan on aggregate %s is complete.

### Parameters

**aggr** (STRING): Aggregate name.

## **wafldb.scan.br.redir.abort**

### Severity

INFORMATIONAL

### Description

This message occurs when a redirect scan is aborted on a volume or aggregate. The redirect scan improves performance following block reallocation by updating metadata that was changed during reallocation. It runs automatically when aggregate reallocation is scheduled using the 'reallocate -A' command.

### Corrective Action

(None).

### Syslog Message

Redirect scan on %s %s%s%s%s has been aborted.

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume or aggregate owner.

**vol** (STRING): Volume or aggregate name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## **wafldb.scan.br.redir.done**

### Severity

INFORMATIONAL

## Description

This message occurs when a redirect scan is complete on a volume or aggregate. The redirect scan improves performance following block reallocation by updating metadata that was changed during reallocation. It runs automatically when aggregate reallocation is scheduled using the 'reallocate -A' command.

## Corrective Action

(None).

## Syslog Message

Redirect scan on %s %s%s%s%s is complete. The scanner took %ld ms.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume or aggregate owner.

**vol** (STRING): Volume or aggregate name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**run\_time** (LONGINT): Total run time of the scanner.

## waf1.scan.br.redir.noStart

### Severity

NOTICE

### Description

This message occurs when a redirect scan cannot be started on a volume or aggregate. The redirect scan improves performance following block reallocation by updating metadata that was changed during reallocation. It runs automatically when aggregate reallocation is scheduled using the 'reallocate -A' command.

### Corrective Action

(None).

### Syslog Message

Redirect scan on %s %s%s%s%s could not be started.

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume or aggregate owner.

**vol** (STRING): Volume or aggregate name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf1.scan.cantStart

### Severity

INFORMATIONAL

### Description

The message occurs when the scan could not start because the maximum number of scans of the specified type are already running.

## Corrective Action

The maximum number of scan operations has been reached. Please wait until the current scan operations have completed before trying any new ones.

## Syslog Message

Cannot start scan %s on %s %s%s%s because maximum number of scans (%d) already running.

## Parameters

**scantype** (STRING): The type of scan

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**max** (INT): The maximum number of scans that can run

## waf.scan.ckpt.init.failed

### Severity

NOTICE

### Description

This message occurs when the creation of Scan Checkpoint Metafile fails.

### Corrective Action

(None).

### Syslog Message

On %s %s%s%s, FSID 0x%x, Creation of Scan Checkpoint Metafile failed with error %d.

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**fsid** (INT): File system identifier of the volume.

**error** (INT): Error number.

## waf.scan.ckpt.remove.failed

### Severity

NOTICE

### Description

This message occurs when the removal of Scan Checkpoint in Metafile fails.

### Corrective Action

(None).

## Syslog Message

On %s %s%s%s%s, FSID 0x%x, Removal of Scan Checkpoint failed with error %d.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**fsid** (INT): File system identifier of the volume.

**error** (INT): Error number.

## wabl.scan.clone.split.abort

### Severity

INFORMATIONAL

### Description

This event is issued to indicate that an operation to split a clone flexible volume from its parent volume was aborted.

### Corrective Action

(None).

## Syslog Message

%s %s%s%s, clone split operation stopped.

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wabl.scan.clone.split.cantLock

### Severity

ERROR

### Description

This event is issued when the system cannot create a volume reference to complete a clone split operation. Check the type to determine if the reference is for the volume or for its enclosing aggregate.

### Corrective Action

Determine why the volume reference failed, and retry the operation if necessary.

## Syslog Message

Unable to lock %s %s%s to complete clone split operation.

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wafI.scan.clone.split.complete**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued to indicate that an operation to split a clone flexible volume from its parent volume completed successfully.

### **Corrective Action**

(None).

### **Syslog Message**

Clone split complete for %s %s%s%s

### **Parameters**

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wafI.scan.clone.split.fail**

### **Severity**

INFORMATIONAL

### **Description**

This event is issued to indicate that an operation to split a clone flexible volume from its parent volume has failed.

### **Corrective Action**

(None).

### **Syslog Message**

%s %s%s%s, clone split operation failed.

### **Parameters**

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wafI.scan.css.undo.complete**

**Severity**

INFORMATIONAL

**Description**

This event is issued to indicate that an operation to undo sharing-by-split in a flexible volume has completed.

**Corrective Action**

(None).

**Syslog Message**

Clone sharing-by-split undo complete for %s %s%s

**Parameters**

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**waf1.scan.deswizzle.aborted****Severity**

ERROR

**Description**

This message occurs when the deswizzler scanner encounters an error and aborts.

**Corrective Action**

Correct the issue based on the reason the deswizzler scanner aborted.

**Syslog Message**

The deswizzler scanner running on volume %s (Snapshot ID %d) was aborted. Reason: "%s".

**Parameters**

**vol** (STRING): Volume that was being deswizzled.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**reason** (STRING): Reason the deswizzler scanner aborted.

**waf1.scan.deswizzle.check.referenced****Severity**

INFORMATIONAL

**Description**

This event is triggered when deswizzling a block already pointed to by pvbn in an older snapshot. If snap1 is not 0 (i.e. not the active fs), there is a strong possibility this may lead to a corruption.

**Corrective Action**

(None).



## Syslog Message

(None).

## Parameters

**vol** (STRING): The volume name  
**snap1** (INT): The snapshot of the block being deswizzled.  
**snap2** (INT): The next older snapshot.  
**cp1** (INT): CP count of snap1.  
**cp2** (INT): CP count of snap2.  
**ino\_type** (STRING): The inode type  
**inum** (INT): Inode containing the block being deswizzled.  
**fbn** (LONGINT): FBN of the block in the inode.  
**level** (INT): Level of the block.  
**vvbn** (LONGINT): VVBN of the block in snap1.  
**pvbn1** (LONGINT): PVBN of the block in snap1.  
**pvbn2** (LONGINT): PVBN of the block in snap2.  
**parent** (LONGINT): VVBN of the parent block in snap2.

## wafI.scan.deswizzle.exception

### Severity

ERROR

### Description

This message occurs when the deswizzler scanner encounters an exception.

### Corrective Action

(None).

## Syslog Message

The deswizzler scanner running on volume %s (Snapshot ID %d) was excepted. Error: "%d".

## Parameters

**vol** (STRING): Volume that was being deswizzled.  
**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy  
**error** (INT): Message status.

## wafI.scan.deswizzle.failed

### Severity

ERROR

### Description

This message occurs when the deswizzler scanner has encounters a volume inconsistency and stops running.

### Corrective Action

An inconsistency was found in the specified Snapshot copy. Delete the Snapshot copy. To ensure that the active file system is intact, you might want to run wafIiron.

## Syslog Message

The deswizzler scanner on flexvol %s, Snapshot copy %s (ID %d) encountered a volume inconsistency and stopped running on the volume.

## Parameters

**vol** (STRING): Volume that was being deswizzled.  
**snapname** (STRING): Inconsistent Snapshot(tm) copy name.  
**snapid** (INT): Inconsistent Snapshot(tm) copy snapid.

## wafl.scan.done

### Severity

NOTICE

### Description

This message occurs when a WAFL® scan is completed.

### Corrective Action

(None).

## Syslog Message

Completed %s on %s %s%s%s%s.

## Parameters

**scantype** (STRING): Type of scan.  
**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## wafl.scan.filecheck.done

### Severity

NOTICE

### Description

This event is issued when the 'file check' diagnostic tool finishes scanning a file for corrupt data blocks.

### Corrective Action

(None).

## Syslog Message

Writing WAFL filecheck log %s

## Parameters

**filename** (STRING): The name of the log file.

## wafl.scan.i2p.cantAbort

**Severity**

INFORMATIONAL

**Description**

This message occurs during an attempt to abort the I2P scan by using the "wafI scan abort" command while running clustered Data ONTAP®. The I2P scan cannot be interrupted.

**Corrective Action**

(None).

**Syslog Message**

You cannot abort an I2P scan on a cluster volume.

**Parameters**

(None).

**wafI.scan.iron.done****Severity**

INFORMATIONAL

**Description**

This message informs the user that the system has completed a wafIiron scan.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s%s%s, wafIiron %s.

**Parameters**

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**status** (STRING): Status string.

**wafI.scan.iron.snapshot****Severity**

INFORMATIONAL

**Description**

This message informs the user that the system has started wafIiron scan on the snapshot.

**Corrective Action**

(None).

**Syslog Message**

%s %s%s%s%s, wafIiron started on snapshot %d.

## Parameters

**type** (STRING): Type of object (volume/aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.  
**snapid** (INT): The snapshot Identifier.

## wafl.scan.iv.metatype.done

### Severity

INFORMATIONAL

### Description

This message occurs when the system completes a WAFL\_FLAG\_IV\_METATYPE inode flag scan on the Infinite Volume namespace constituent.

### Corrective Action

(None).

### Syslog Message

The WAFL\_FLAG\_IV\_METATYPE inode flag scanner has completed the scan on %s %s%s%s%s. The scanner scanned %d inodes, tagged %d inodes with the inode flag, and took %ld msec to complete the scan.

## Parameters

**type** (STRING): Type of object (volume).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.  
**inodes\_visited** (INT): Total number of inodes visited by the scanner.  
**inodes\_tagged** (INT): Total number of inodes tagged with the WAFL\_FLAG\_IV\_METATYPE inode flag by the scanner.  
**run\_time** (LONGINT): Total run time of the scanner in milliseconds.

## wafl.scan.iv.metatype.dud

### Severity

ERROR

### Description

This message occurs when the system tries and fails to start the WAFL\_FLAG\_IV\_METATYPE inode flag scan on the Infinite Volume namespace constituent. This might cause errors in file operations over NFS/CIFS in the Infinite Volume until the scanner is restarted and the scan is completed.

### Corrective Action

The scanner will restart automatically in cases where the failure to start occurs due to a storage failover event, an in-progress node revert operation, or a node reboot. In other cases, contact NetApp technical support for assistance.

## Syslog Message

Failed to launch the WAFL\_FLAG\_IV\_METATYPE inode flag scanner on %s %s%s%s%s. The error code returned was %d (%s).

## Parameters

**type** (STRING): Type of object (volume).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**err\_code** (INT): Numerical error code indicating the reason the scan could not be launched.

**err\_string** (STRING): String interpreting the error code.

## wافل.scan.iv.metatype.start

### Severity

INFORMATIONAL

### Description

This message occurs when the system starts a WAFL\_FLAG\_IV\_METATYPE inode flag scan on the Infinite Volume namespace constituent.

### Corrective Action

(None).

## Syslog Message

The WAFL\_FLAG\_IV\_METATYPE inode flag scanner has started running on %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## wافل.scan.l1.uncompact.done

### Severity

NOTICE

### Description

This message occurs when the operation to uncompact indirect blocks in a FlexVol® volume is complete.

### Corrective Action

(None).

## Syslog Message

L1 uncompact operation complete for %s (UUID: %s). %d indirect blocks were uncompactd.

## Parameters

**name** (STRING): Volume name.

**vol\_uuid** (STRING): Volume UUID.  
**count** (INT): Number of indirect blocks being uncompactd.

## **waf1.scan.l1.uncompact.nobt**

### **Severity**

ERROR

### **Description**

This message occurs when an indirect block uncompact operation is attempted on a volume before block type metadata initialization is complete.

### **Corrective Action**

Check whether the block type scan is still running on the volume. Try the indirect block uncompact operation on the volume after the block type scan is complete.

### **Syslog Message**

Indirect block uncompact scan can't be started on volume %s (UUID: %s) because block type initialization is not complete on the volume. Try again after block type scan is complete on the volume.

### **Parameters**

**name** (STRING): Volume name.  
**vol\_uuid** (STRING): Volume UUID.

## **waf1.scan.l1.uncompact.nospc**

### **Severity**

ERROR

### **Description**

This message occurs when there is not enough free space to uncompact the indirect blocks during the volume uncompact operation.

### **Corrective Action**

Create space by increasing the volume size, or by deleting data or deleting Snapshot(tm) copies. To increase a volume's size, use the "volume size" command. To delete a volume's Snapshot copies, use the "volume snapshot delete" command.

### **Syslog Message**

Indirect block uncompact scan for volume %s (UUID: %s) stopped because of insufficient free space.

### **Parameters**

**name** (STRING): Volume name.  
**vol\_uuid** (STRING): Volume UUID.

## **waf1.scan.layout.advise**

### **Severity**

INFORMATIONAL

## Description

This event is issued to advise the operator of optimal settings for their layout ratio.

## Corrective Action

(None).

## Syslog Message

WAFS layout ratio for %s %s%s%s%s is %d.%02d. A ratio of 1 is optimal. Based on your free space, %d.%02d is expected.

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**ratio1** (INT): The ratio integer

**ratio2** (INT): The ratio fraction of 100

**eratio1** (INT): The expected ratio integer

**eratio2** (INT): The expected ratio fraction of 100

## wafL.scan.layout.advise.ino

### Severity

INFORMATIONAL

### Description

This event is issued to advise the operator of optimal settings for their layout ratio.

### Corrective Action

(None).

### Syslog Message

WAFS layout ratio for %s %s%s%s%s, %s inode %u is %d.%02d. A ratio of 1 is optimal. Based on your free space, %d.%02d is expected.

### Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**inotype** (STRING): The inode space (public or private)

**ino** (INT): The inode value

**ratio1** (INT): The ratio integer

**ratio2** (INT): The ratio fraction of 100

**eratio1** (INT): The expected ratio integer

**eratio2** (INT): The expected ration fraction of 100

## wafL.scan.layout.cantWrite

**Severity**

ERROR

**Description**

This message occurs when the system cannot write out the WAFL®scan layout log.

**Corrective Action**

The most likely cause of a failure in writing the file is lack of space on the specified volume. Use the 'df' command to check available space. If the volume is full, free some space on the volume, and then restart the command.

**Syslog Message**

Cannot write WAFL scan layout log %s.

**Parameters**

**filename** (STRING): Name of the log file.

**wافل.scan.layout.write****Severity**

NOTICE

**Description**

This event is issued when the system initiate writing of the scan layout log.

**Corrective Action**

(None).

**Syslog Message**

Writing WAFL layout log %s

**Parameters**

**filename** (STRING): The name of the log file

**wافل.scan.log.cantWrite****Severity**

NOTICE

**Description**

This message occurs when the system cannot write out the scan report log.

**Corrective Action**

If the error is due to lack of space, add more space or free space in the volume. If the problem persists, contact NetApp technical support.

**Syslog Message**

Cannot write WAFL vbntoino log %s: %d (%s).

**Parameters**

**filename** (STRING): Name of the log file.



**err\_code** (INT): A numerical error code indicating the reason why the log was not written.  
**err\_string** (STRING): A string interpreting the error code.

## **waf1.scan.lost.obj.recovery**

### **Severity**

NOTICE

### **Description**

This message occurs after a lost object recovery scanner has been completed or aborted.

### **Corrective Action**

None.

### **Syslog Message**

Lost object recovery scanner on volume "%s%s%s" is %s. (num\_lost=%llu, num\_inserted=%llu, num\_deleted=%llu, num\_skipped=%llu, num\_objs\_in\_namespace=%llu)

### **Parameters**

**volname** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**state** (STRING): State of the scan, either aborted or completed

**num\_lost** (LONGINT): Number of objects found that were not in the namespace.

**num\_inserted** (LONGINT): Number of objects fixed by adding them back into the namespace.

**num\_deleted** (LONGINT): Number of objects fixed by deletion.

**num\_skipped** (LONGINT): Number of inodes skipped during scanning due to some error.

**num\_objs\_in\_namespace** (LONGINT): Number of objects that were found to exist already in the namespace.

## **waf1.scan.ownblocks.done**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a scan operation to check owned blocks is complete.

### **Corrective Action**

(None).

### **Syslog Message**

Completed %s on %s %s%s%s%s. The scanner took %ld ms.

### **Parameters**

**op** (STRING): Name of the operation.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**run\_time** (LONGINT): Total run time of the scanner.

## **wafl.scan.quota.init.cantAbort**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs during an attempt to abort the quota initialization scan by using the "wafl scan abort" command. The quota initialization scan cannot be interrupted.

### **Corrective Action**

Wait for the quota initialization to complete or use the "quota off" command.

### **Syslog Message**

User cannot abort quota initialization scan. Please use 'quota off' command.

### **Parameters**

(None).

## **wafl.scan.realloc.oldsnap**

### **Severity**

ERROR

### **Description**

This message occurs when a block reallocation scanner cannot move blocks trapped in aggregate Snapshot® copies, because the Snapshot copies are older than the value in 'wafl\_blk\_max\_snap\_age' flag (default is 360 seconds). You can display the value of this flag by using 'printflag wafl\_blk\_max\_snap\_age'.

### **Corrective Action**

Delete all aggregate Snapshot copies or Snapshot schedules before running a block reallocation scan. In the case of RaidSyncmirror aggregates, ensure that the resyncsnaptime option has a value that is less than in the 'age' parameter.

### **Syslog Message**

Block reallocation scan on aggregate %s found snapshots older than %d seconds.

### **Parameters**

**aggr** (STRING): Aggregate name.

**age** (INT): Maximum allowed age of aggregate Snapshot copies, (in seconds), for block reallocation scan.

## **wafl.scan.reallocDisallowed**

### **Severity**

ERROR

### **Description**

This event is issued when an attempt is made to reallocate a file or LUN that cannot be reallocated. Examples of such files/LUNs include file in snapshots, or Qtree Snapmirror destinations.

### **Corrective Action**

File/LUN reallocation should be specified on files/LUNs in the active filesystem that are not part of a QSM destination.

## Syslog Message

Reallocation disallowed on %s inode %d in volume %s%s: readonly or snapshot.

## Parameters

**inotype** (STRING): The inode space (public or private)

**inode** (INT): The inode that could not be reallocated.

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.scan.relocation.nospaceFree

### Severity

ERROR

### Description

This event is issued to indicate there is not enough free space to rewrite the blocks during file or volume reallocation.

### Corrective Action

Free space in the volume, or increase the size of a flexible volume, then retry the operation.

## Syslog Message

Reallocate scan for volume %s%s has stopped because there is insufficient free space.

## Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.scan.relocation.nospaceSnap

### Severity

ERROR

### Description

This event is issued to indicate there is not enough snapshot free space to rewrite the blocks during file or volume reallocation.

### Corrective Action

Remove snapshots in the volume, or increase the size of a flexible volume, then retry the operation.

## Syslog Message

Reallocate scan for volume %s%s has stopped because there is insufficient snapshot free space.

## Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## **wافل.scan.snap.vbnOutOfRange**

### **Severity**

ERROR

### **Description**

This message occurs when scanners find an out-of-range block number in the Snapshot(tm) copy.

### **Corrective Action**

Delete the specified Snapshot copy.

### **Syslog Message**

Out of range block number error detected at %s %s%%s%%s in Snapshot copy %d, during block type scan.  
Delete the specified Snapshot copy.

### **Parameters**

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**snapId** (INT): Identifier for the Snapshot copy.

## **wافل.scan.split.inofile.cantAbort**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs during an attempt to abort the split inofile scan by using the "wافل scan abort" command. The split inofile scan cannot be interrupted.

### **Corrective Action**

(None).

### **Syslog Message**

User cannot abort the split inofile scan.

### **Parameters**

(None).

## **wافل.scan.typebits.diffFail**

### **Severity**

ALERT

### **Description**

This message occurs when the type bit scan encounters an unexpected error from the snapdiff APIs. This can imply that prev\_snapId snapshot might be corrupt.

## Corrective Action

Contact NetApp technical support for corrective steps.

## Syslog Message

Aborting the type bit scan due to error %d on volume %s while running a diff operation on inode %u of %s space between Snapshot ID %u and Snapshot ID %u after fbn %llu.

## Parameters

**error** (INT): Error number.  
**volname** (STRING): Volume name.  
**inum** (INT): File ID under scan.  
**ino\_type** (STRING): Inode type.  
**snapId** (INT): Identifier for the Snapshot(tm) copy under scan.  
**prev\_snapId** (INT): Identifier for the reference Snapshot copy.  
**start\_fbn** (LONGINT): Starting file block number (fbn) of the diff operation.  
**vendor** (STRING): Name of the vendor.

## wafI.scan.typeBitsMismatch

### Severity

NOTICE

### Description

This messages occurs when scanners find a type bits mismatch for a block number. This might indicate a file system inconsistency in a file system or Snapshot(tm) copy.

### Corrective Action

(Call support).

## Syslog Message

Type bits mismatch found at %s %s%s%s%s in Snapshot copy %d, vbn %llu, during block type scan. On-disk type bits %d, expected type bits %d.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Node that owns the object.  
**vol** (STRING): Name of the object.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the object in cases where the object name itself is insufficient.  
**snapId** (INT): Identifier for the Snapshot copy.  
**vbn** (LONGINT): Block number of the block within the volume.  
**diskbits** (INT): On-disk type bits of the block.  
**expectedbits** (INT): Expected type bits of the block.

## wafI.scan.uncompact.aborted

### Severity

ERROR

### Description

This message occurs when the uncompression scanner is aborted.

## Corrective Action

(None).

## Syslog Message

%s scanner aborted on %s %s%s%s%s with error: %s. %s was tried on %ld blocks before failure.

## Parameters

**opname** (STRING): Opname for the scanner.  
**voltype** (STRING): Volume type on which scanner was run.  
**aggr\_owner** (STRING): Owner of the volume.  
**aggr\_name** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**uuid** (STRING): Volume uuid.  
**abort\_reason** (STRING): Reason for stopping scan.  
**opname1** (STRING): Opname for the scanner.  
**val** (LONGINT): Number of vbns acted on.

## wafl.scan.vbn.outOfRange

### Severity

NOTICE

### Description

This message occurs when scanners find an out-of-range block number in the active file system.

### Corrective Action

(None).

## Syslog Message

Out-of-range block number error detected at %s %s%s%s%s, during block type scan.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafl.scan.wafliron.cantAbort

### Severity

INFORMATIONAL

### Description

This message occurs during an attempt to abort the wafliron scan by using the "wafl scan abort" command. The wafliron scan cannot be interrupted.

### Corrective Action

Wait for the wafliron scan to finish or use "aggr wafliron stop" to stop the scan.

## Syslog Message

User cannot abort wafiron scan. Please use 'vol wafiron stop' command.

## Parameters

(None).

# wafl.seclabel events

## wafl.seclabel.syncsm.mutex

### Severity

NOTICE

### Description

This message occurs when the mutual exclusion has failed between the SnapMirror Synchronous and NFS v4.2 Security Labels features. This happens when an existing volume participating in a SnapMirror relationship receives an NFS "open", "create", or "setattr" operation carrying a security label. The system will not let a security label operation continue on a volume that participates in a SnapMirror Synchronous relationship, so the system administrator must either stop mirroring the volume or drop the security label operation completely.

### Corrective Action

(None).

## Syslog Message

An incompatible NFS v4.2 security label operation was received on a file in volume "%s" having a SnapMirror Synchronous relationship. To allow NFS operations carrying security label on the volume, the mirroring must be stopped first: FSID "%d" SNAPID "%d" FILEID "%llu": "%s"

## Parameters

**volname** (STRING): Volume name.

**fsid** (INT): File System ID (FSID) of the volume.

**snapid** (INT): Snapshot copy ID.

**fileid** (LONGINT): File ID.

**error\_msg** (STRING): Detailed error message to provide context while debugging the issue. The relevant input/output parameters of the message are also included to aid in debugging.

# wafl.sfsr events

## wafl.sfsr.abort.ino

### Severity

NOTICE

### Description

This event is issued when the system is forced to abort a file restore.

### Corrective Action

(None).

## Syslog Message

Single-file snaprestore of inode %u (snapid %d, %s %s%s%s%s) to inode %u prematurely halted.

## Parameters

**srcIno** (INT): The source inode number

**snapId** (INT): The snapshot identifier

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**dstIno** (INT): The target inode number

## waf.sfsr.abort.inoDeleted

### Severity

NOTICE

### Description

This event is issued when the system is forced to abort a file restore. The inode was deleted.

### Corrective Action

(None).

## Syslog Message

Single-file snaprestore of inode %u (snapid %d, %s %s%s%s%s) to inode %u aborted. File was deleted.

## Parameters

**srcIno** (INT): The source inode number

**snapId** (INT): The snapshot identifier

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**dstIno** (INT): The target inode number

## waf.sfsr.done

### Severity

INFORMATIONAL

### Description

This message occurs when a file restore operation has completed.

### Corrective Action

(None).

## Syslog Message

Single-file snaprestore of inode %u (snapid %d, %s %s%s%s%s) to inode %u has completed.



## Parameters

**srcIno** (INT): Source inode number.  
**snapId** (INT): Snapshot™ identifier.  
**type** (STRING): Type of object (volume/aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.  
**dstIno** (INT): Target inode number.

## waf.sfsr.lun.error

### Severity

ERROR

### Description

This message occurs when the system attempts to restore a file onto a LUN. This is not a supported action; files cannot be restored onto a LUN.

### Corrective Action

No corrective action is possible; files cannot be restored onto a LUN.

### Syslog Message

Attempted to restore a file onto a LUN, but files cannot be restored onto a LUN. (vol='%s', volident='%s', snapname='%s', filename='%s')

## Parameters

**vol** (STRING): Volume name.  
**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.  
**snapname** (STRING): Name of the Snapshot(tm) copy that the file is being restored from.  
**filename** (STRING): Name of the file that cannot be restored onto a LUN.

## waf.sfsr.lun.name.error

### Severity

ERROR

### Description

This event is issued when a LUN is being restored to a new location and the path to the new location contains invalid characters. The valid characters for the last component in the lun path are A-Z, a-z, 0-9, '-', ':', '\_', '{' and '}'.

### Corrective Action

Make sure that the last component of the path to the new location contains characters from A-Z, a-z, 0-9, '-', ':', '\_', '{' and '}'.

### Syslog Message

Can't restore LUN due to invalid restore path. The last component of the restore path can only contain characters from A-Z, a-z, 0-9, '-', ':', '\_', '{' and '}'.

## Parameters

(None).

## waf1.sfsr.lun.restore.error

### Severity

ERROR

### Description

This message occurs when the system attempts to restore a LUN to something other than qtree root. This is not a supported action; LUNs can only be restored into qtree roots.

### Corrective Action

No corrective action is possible; LUNs can only be restored into qtree roots.

### Syslog Message

Attempted to restore a LUN to something other than a qtree root, but LUNs can only be restored into qtree roots. (vol='%s', volident='%s', snapname='%s', filename='%s')

### Parameters

**vol** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**snapname** (STRING): Name of the Snapshot(tm) copy that the LUN is being restored from.

**filename** (STRING): Name of the file that the LUN is being restored from.

## waf1.sfsr.sr.disabled

### Severity

ERROR

### Description

This event is generated when a single-file snap restore (sfsr) operation on a file with a space reservation fails because there is not enough space on the volume to restore the file with its space reservation. If there was room for the file without a space reservation, then the file was restored with its space reservation disabled.

### Corrective Action

If this event happened because the single-file snaprestore (sfsr) failed then space on the volume must be freed and the sfsr redone. If the sfsr was otherwise successful then space on the volume must be freed and reservations re-enabled by using "file reservation file-path enable".

### Syslog Message

Single-file snaprestore of inode %u on %s %s%s%s) did not restore space reservations due to lack of disk space.

### Parameters

**srcino** (INT): The source inode number

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# wafl.shutdown events

## wafl.shutdown.notClean

### Severity

NOTICE

### Description

This message occurs during shutdown, to indicate that a volume has not been marked clean. By itself, this is not necessarily a problem. When this event occurs in conjunction with WAFL inconsistency panics or higher-severity EMS events, this message can indicate a potential source of the inconsistencies.

### Corrective Action

No immediate action is required. If this message occurs in conjunction with WAFL inconsistency panics or other EMS messages, contact NetApp technical support.

### Syslog Message

%s %s%s%s%s not marked clean (volume state is %s).

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**state** (STRING): State of the volume.

# wafl.sidl events

## wafl.sidl.blkincons.recovery

### Severity

NOTICE

### Description

This message occurs when a volume or aggregate with single instance data logging enabled encounters a problem in a metafile required for the functioning of the single instance data logging feature. The system will attempt automatic recovery within the next couple of Consistency Points (CPs). During the period of recovery, write performance could be lower than usual.

### Corrective Action

(None).

### Syslog Message

%s:%s encountered a problem in a metafile (sidl). The system will automatically repair the problem.

### Parameters

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).

**object\_name** (STRING): Name of the object.

**fileid** (INT): File identifier of the metafile.

**IO\_blk\_cnt** (LONGINT): Number of L0 blocks present.

**exp\_blk\_cnt** (LONGINT): Number of L0 blocks expected.  
**cpcount** (INT): CP that performs automatic recovery.

## **waf1.sid1.l0incons.recovery**

### **Severity**

NOTICE

### **Description**

This message occurs when a volume or aggregate with single instance data logging enabled encounters a problem in a metafile required for the functioning of the single instance data logging feature. The system will attempt automatic recovery within the next couple of Consistency Points (CPs). During the period of recovery, write performance could be lower than usual.

### **Corrective Action**

(None).

### **Syslog Message**

%s:%s encountered a problem in a metafile (sid1). The system will automatically repair the problem.

### **Parameters**

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).  
**object\_name** (STRING): Name of the object.  
**fbn** (LONGINT): File block number (fbn) where the problem was encountered.  
**fileid** (INT): File identifier of the metafile.  
**vbn** (LONGINT): Value of the volume block number (vbn) that had a problem.  
**where** (STRING): The module that generated the event (write or refill).  
**cpcount** (INT): CP that performs automatic recovery.

## **waf1.sid1.space.disable**

### **Severity**

NOTICE

### **Description**

This message occurs when a volume or aggregate disables the single instance data logging feature because the volume or aggregate is too close to full. During the period when single instance data logging is disabled, write performance could be lower than usual.

### **Corrective Action**

Decrease the space utilization in the volume or aggregate either by deleting data or by growing the volume or aggregate.

### **Syslog Message**

%s:%s has disabled the single instance data logging feature. Write performance could be lower than usual.

### **Parameters**

**object\_type** (STRING): Identifier for the type of object to which this event applies (aggregate or volume).  
**object\_name** (STRING): Name of the object.  
**cpcount** (INT): Consistency Point (CP) at which the event is generated.

# waf.sis events

## waf.sis.clone.autodelete.file

### Severity

NOTICE

### Description

This message occurs when a file clone is autodeleted

### Corrective Action

(None).

### Syslog Message

'%s' is sent for auto deletion.

### Parameters

**path** (STRING): File path.

## waf.sis.clone.autodelete.threshold

### Severity

NOTICE

### Description

This message occurs when file clone create operation fails because volume crossed autodelete threshold.

### Corrective Action

Increase the space available to the volume by either adding disks or by freeing space in the volume to allow the cloning operation to continue.

### Syslog Message

File clone creation cannot be completed because volume crossed autodelete threshold.

### Parameters

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

## waf.sis.clone.backup.file.cleanup.failed

### Severity

ERROR

### Description

This message occurs when clone create gets an error while cleaning up the backup file of the pre-existing clone.

### Corrective Action

A backup of the pre-existing clone file is saved at the backup file location. Its cleanup needs manual intervention.

## Syslog Message

File clone create for the source %s and clone %s returned error while cleaning up the backup file %s of the pre-existing clone with error: %d. Backup file should be deleted to reclaim space.

## Parameters

**src\_filename** (STRING): Source filename.  
**clone\_filename** (STRING): Clone filename.  
**backup\_filename** (STRING): Backup filename.  
**error** (INT): Error returned from FlexClone® subsystem.

## waf.sis.clone.backup.file.restore.failed

### Severity

ERROR

### Description

This message occurs when clone create errors out while restoring the backup file to the pre-existing clone.

### Corrective Action

Clone create failed after deleting the pre-existing clone. A backup of the pre-existing clone file is saved at the backup file location. It can be used to recover the pre-existing clone file.

## Syslog Message

File clone create for the source %s and clone %s failed while restoring the backup file %s to the pre-existing clone with error: %d. Backup file can be moved to the pre-existing clone location for recovering the clone file.

## Parameters

**src\_filename** (STRING): Source filename.  
**clone\_filename** (STRING): Clone filename.  
**backup\_filename** (STRING): Backup filename.  
**error** (INT): Error returned from FlexClone® subsystem.

## waf.sis.clone.create.failed

### Severity

INFORMATIONAL

### Description

This message occurs when attempts to create a file clone fail.

### Corrective Action

(None).

## Syslog Message

File clone create of clone file %s (fileid: %d) from source file %s (fileid: %d) in volume %s failed with error:%s at line %d with src\_fsid %d, dst\_fsid %d, src\_snapid %d, dir\_fileid %d, type %s, category %s.

## Parameters

**clone\_filename** (STRING): Clone filename.  
**dst\_fileid** (INT): File identifier for destination file.

**src\_filename** (STRING): Source filename.  
**src\_fileid** (INT): File identifier for source file.  
**volname** (STRING): Volume name.  
**error** (STRING): Error returned from the FlexClone® subsystem.  
**line\_number** (INT): Line number where the error has occurred.  
**src\_fsid** (INT): File system identifier of source volume.  
**dst\_fsid** (INT): File system identifier of destination volume.  
**snapid** (INT): Snapshot copy identifier.  
**dir\_fileid** (INT): File identifier for the directory.  
**type** (STRING): Type of Cloning.  
**category** (STRING): Clone create requestor information.

## **waf1.sis.clone.unlock.snapshot**

### **Severity**

NOTICE

### **Description**

This event is issued when the volume options are set such that the autodelete destroy policy is allowed to destroy backing snapshots for sis clones. The snapshot autodelete destroy policy is used as a last resort to recover space in a thin provisioned aggregate. The disruptive effect of destroying sis clone backing snapshots should be understood when setting this policy.

### **Corrective Action**

(None).

### **Syslog Message**

Unlocked sis clone backing snapshot in volume %s. Clone files based on snapshot (%d, %d) might become inconsistent.

### **Parameters**

**name** (STRING): The volume name  
**snapid** (INT): Snapshot(tm) copy identifier.  
**cp\_cnt** (INT): The CP(Consistency Point) count. The (snapid, cp\_count) pair uniquely identifies a Snapshot copy.

## **waf1.slice events**

### **waf1.slice.invalid.stream**

#### **Severity**

NOTICE

#### **Description**

This message occurs when an invalid slice stream was not fixed by waf1iron because the associated slice base file was not recognized at the time waf1iron ran. Once this invalid stream file is removed, 'waf1 slice\_repair' can recreate and repair the stream file.

#### **Corrective Action**

A slice stream file in the slicevol is not valid. Please remove the file and rerun 'waf1 slice\_repair' to create a valid slice stream file - Use the (privilege :diag) "rm -x [stream\_file\_name] [slice\_base\_path]" command to remove the invalid slice stream file in the slicevol. Then use the (privilege:diag) "waf1

slice\_repair[slice\_base\_path] start" command to repair the stream file. If the problem persists, contact NetApp technical support.

### Syslog Message

Invalid slice stream file: %s (fileid %d) in the slice base path %s in volume %s%s%s detected. Please run the following commands in the diag mode: %s to fix it.

### Parameters

**stream\_name** (STRING): The name of an invalid slice stream file.  
**stream\_fileid** (INT): The file ID of an invalid slice stream file.  
**base\_path** (STRING): The slice base file path that is being repaired.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Vserver's Universal Unique Identifier (UUID).  
**repair\_cmd** (STRING): The repair command to restore the stream file to be valid.

## waf.sm events

### waf.sm.vvol.read.failed

#### Severity

ALERT

#### Description

This message occurs when SnapMirror® fails to read a block from the container file of a FlexVol® SnapMirror source volume. This error typically means that there is some corruption in the FlexVol volume.

#### Corrective Action

Contact NetApp technical support for assistance.

#### Syslog Message

SnapMirror encountered an error while reading the block at VVBN %llu in volume '%s%s%s'.

#### Parameters

**vvbn** (LONGINT): Virtual volume block number (VVBN) on which the read failed.  
**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Uniquely identifies the volume, if the volume name itself is insufficient.

## waf.snap events

### waf.snap.autoDel.policyChg

#### Severity

NOTICE

#### Description

This message occurs when the system detects a possible change in the autodelete or autosize policy of the volume. This can happen when a version of Data ONTAP® is loaded that might not support the reclaim policy currently set in the active filesystem.



## Corrective Action

Unrecognized reclaim policy settings are cleared in the active file system. If you subsequently load a version of Data ONTAP® that recognizes the reclaim policy settings, you need to reestablish the policy settings using 'snap autodelete' or 'vol autosize' commands.

## Syslog Message

Reclaim policy might have changed in %s %s%s%s%s. Check the 'snap autodelete' or 'vol autosize' command settings.

## Parameters

**type** (STRING): Container type, either volume or aggregate.

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## wapl.snap.autoDelete

### Severity

INFORMATIONAL

### Description

This message occurs when the aggregate Snapshot™ copy reserve is full and Snapshot copies are being deleted to free space.

### Corrective Action

If aggregate Snapshot copies are being deleted too frequently or while they are still being used, then consider the following options: o Increase the Snapshot copy reserve percentage for the aggregate. o Alter the aggregate Snapshot copy schedule to create less frequent Snapshot copies.

### Syslog Message

Deleting Snapshot copy '%s' in aggregate '%s' to free up space.

### Parameters

**name** (STRING): Snapshot copy name.

**aggr** (STRING): Aggregate name.

## wapl.snap.autoDelete.createStateSnap.fail

### Severity

INFORMATIONAL

### Description

This event indicates that after auto-deletion of all aggregate snapshots, we have failed to take a state snapshot.

### Corrective Action

(None).

### Syslog Message

Unable to create snapshot '%s' in aggregate '%s': %s

## Parameters

**name** (STRING): The snapshot name  
**vol** (STRING): The volume name  
**err** (STRING): Reason why the operation failed.

## waf.snap.autoDelete.fail

### Severity

NOTICE

### Description

This event indicates that the automatic deletion of a snapshot in an aggregate has failed. An appliance may automatically delete an aggregate snapshot if the aggregate snap reserve fills.

### Corrective Action

The appliance will automatically retry the snapshot deletion shortly after the failure occurs. If the deletion succeeds then no corrective action is necessary. If deletion repeatedly fails then the aggregate could potentially become full. To address this please contact customer support.

### Syslog Message

Unable to delete snapshot '%s' in aggregate '%s' to recover space: %s

## Parameters

**name** (STRING): The snapshot name  
**aggr** (STRING): The aggregate name  
**err** (STRING): The reason why the deletion failed.

## waf.snap.create

### Severity

NOTICE

### Description

This message occurs when a Snapshot™ copy is created successfully.

### Corrective Action

(None).

### Syslog Message

Snapshot copy %s was created successfully for volume %s%s%s in %ld milliseconds.

## Parameters

**name** (STRING): Snapshot copy name.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.  
**time** (LONGINT): Time taken to create the Snapshot copy, in milliseconds.

## waf.snap.create.skip.reason

## Severity

ERROR

## Description

This message occurs when the creation of a Snapshot(tm) copy is skipped.

## Corrective Action

Address the issue identified by the "reason" parameter. For example, if the reason is "Volume is too full", increase the volume size or free up disk space on the target volume.

## Syslog Message

%s %s %s: Skipping creation of %s Snapshot copy (%s).

## Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier.

**snapname** (STRING): Name of the Snapshot copy being created.

**reason** (STRING): Reason for skipping the Snapshot copy.

## wapl.snap.delete

## Severity

INFORMATIONAL

## Description

This message occurs when a Snapshot™ copy is deleted.

## Corrective Action

(None).

## Syslog Message

Snapshot copy %s on %s %s%s%s%s %s was deleted by the Data ONTAP function %s. The unique ID for this Snapshot copy is (%d, %d).

## Parameters

**snap\_name** (STRING): Name of the deleted Snapshot copy.

**type** (STRING): Container type (volume or aggregate).

**owner** (STRING): Volume or aggregate owner.

**vol** (STRING): Volume or aggregate name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume or aggregate when the volume or aggregate name by itself is insufficient.

**vendor** (STRING): Vendor name.

**caller\_func** (STRING): Function in Data ONTAP® that invokes the deletion process.

**snapid** (INT): Numerical ID of the Snapshot copy. The (snapid, cp\_count) pair uniquely identifies a Snapshot copy.

**cp\_count** (INT): The CP (consistency point) count. The (snapid, cp\_count) pair uniquely identifies a Snapshot copy.

## waf1.snap.delete.32bit

### Severity

INFORMATIONAL

### Description

This message occurs when a Snapshot(tm) copy on an aggregate or volume is deleted because it has an unsupported 32-bit or mixed format file system.

### Corrective Action

(None).

### Syslog Message

A Snapshot copy on %s %s%s%s%s was deleted because it has an unsupported 32-bit or mixed format file system. The unique ID for this Snapshot copy is (%d, %d).

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

**snapid** (INT): Numerical ID of the Snapshot copy. As a pair, the values of the "snapid" and "cp\_count" parameters uniquely identify a Snapshot copy.

**cp\_count** (INT): CP (consistency point) count. As a pair, the values of the "snapid" and "cp\_count" parameters uniquely identify a Snapshot copy.

## waf1.snap.delete.pending

### Severity

INFORMATIONAL

### Description

This message occurs when the deletion of a Snapshot(tm) copy is pending.

### Corrective Action

(None).

### Syslog Message

Pending Snapshot copy (ID %u) deletion in %s %s%s%s%s.

### Parameters

**snapid** (INT): Snapshot copy ID.

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.snap.deleteError

**Severity**

INFORMATIONAL

**Description**

This message occurs when attempts to delete a Snapshot® copy fail.

**Corrective Action**

(None).

**Syslog Message**

Unable to remove old Snapshot copy "%s" on %s %s%s%s. Error: %s.

**Parameters**

**name** (STRING): Name of the Snapshot copy.

**type** (STRING): Type of object (volume or aggregate).

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): UUID of the volume, to uniquely identify the volume when the volume name alone is insufficient.

**err** (STRING): Error description

**waf.snap.sched.notok****Severity**

NOTICE

**Description**

This message occurs when the creation of a Snapshot(tm) copy is skipped:

**Corrective Action**

(None).

**Syslog Message**

%s%s%s%s: %s

**Parameters**

**message** (STRING): Generic scheduled Snapshot copy failure indication.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**reason** (STRING): Specific reason for the Snapshot failure.

**waf.snaplock events****waf.snaplock.scan.complete****Severity**

NOTICE

## Description

This event indicates that the snaplock scan has completed.

## Corrective Action

(None).

## Syslog Message

%s %s%s%s%s SnapLock scan determined upgraded volume expiry date: %s

## Parameters

**typestr\_cap** (STRING): wafL\_Volume type, capitalized

**owner** (STRING): owner string (foreign, partner, et al)

**vol\_name** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**timebuf** (STRING): Text string describing expiration time/date.

# wafL.snaprestore events

## wafL.snaprestore.badVersion

### Severity

ERROR

### Description

This message occurs during a Snapshot(tm) copy restore operation, when the system cannot recognize the version of Data ONTAP® from the Snapshot copy it is reverting to. The Snapshot copy is corrupt and cannot be used.

### Corrective Action

Cancel the Snapshot copy restore operation, if necessary, and restore to a different Snapshot copy.

### Syslog Message

Snapshot copy is from an unrecognizable version of Data ONTAP. (vol='%s', volident='%s')

### Parameters

**vol** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafL.snaprestore.cancel.invalid

### Severity

NOTICE

### Description

This message occurs during volume initialization when the system cancels a SnapRestore® process due to an invalid Snapshot™ copy.

### Corrective Action

Select a different Snapshot copy and reattempt the restore operation.

## Syslog Message

Snapshot is invalid; canceling pending "vol snaprestore" for %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## wافل.snaprestore.cancel.prevcp

### Severity

INFORMATIONAL

### Description

This message occurs during volume initialization from a previous consistency point (CP) when the system cancels a SnapRestore® process. This happens when booting to a previous CP.

### Corrective Action

Retry the SnapRestore operation after the system completes bootup.

## Syslog Message

Canceling pending "vol snaprestore" for %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## wافل.snaprestore.defer.prevcp

### Severity

INFORMATIONAL

### Description

This message occurs during a read-only boot to the previous consistency point (CP). In this case, the pending SnapRestore® process is deferred.

### Corrective Action

(None).

## Syslog Message

Readonly mode; deferring "vol snaprestore" for %s %s%s%s%s.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## **waf.snapstore.revert**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs during volume initialization when the system reverts to a previous Snapshot™ copy.

### **Corrective Action**

(None).

### **Syslog Message**

Reverting %s %s%s%s to a previous Snapshot copy with snapid %d.

### **Parameters**

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**snapid** (INT): Identifier for the Snapshot copy.

## **waf.snapstore.stale**

### **Severity**

ERROR

### **Description**

This message occurs during a Snapshot(tm) copy restore operation, when the system determines that the Snapshot copy it is restoring from is stale. The Snapshot copy is corrupt and cannot be used.

### **Corrective Action**

Cancel the Snapshot copy restore operation, if necessary, and restore from a different Snapshot copy.

### **Syslog Message**

Snapshot copy is stale. (vol='%s', volident='%s')

### **Parameters**

**vol** (STRING): Volume name.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## **waf.snapshot events**

### **waf.snapshot.badFileSize**

#### **Severity**

NOTICE



## Description

This event is issued when the system discover that a filesize is out of range in a snapshot.

## Corrective Action

(None).

## Syslog Message

File size out of range in snapshot.

## Parameters

(None).

## wapl.snapshot.badLevel

### Severity

NOTICE

### Description

This event is issued when there is an inconsistency in the metadata of an inode in a snapshot. This message is advisory only. The inconsistency will be corrected in memory only, and the data will be made available.

### Corrective Action

No action is required; however, the message may repeat until the snapshot is deleted.

### Syslog Message

Corrected bad level %d for buftree of %s fileid %d, snapid %d in %s %s%s%s%s

### Parameters

**level** (INT): The level of the buftree of the file

**ino\_type** (STRING): The inode type

**fileid** (INT): The file identifier

**snapid** (INT): The snapshot identifier

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wapl.snapshot.ecv.fail

### Severity

ERROR

### Description

This message occurs when the system attempts to create a Snapshot(tm) copy on a mixed-version cluster. However, the effective version of the cluster is ONTAP® 9.3 or earlier, and the maximum number of Snapshot copies supported by this effective version has already been reached.

### Corrective Action

Upgrade all nodes in the cluster to ONTAP 9.4 or later.

## Syslog Message

Snapshot copy create operation failed on volume %s%s%s. Maximum number of Snapshot copies reached on a mixed-version cluster.

## Parameters

**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): ID of the volume.

# waf.snapshot\_xlation\_duplicate\_entry events

## waf.snapshot\_xlation\_duplicate\_entry

### Severity

ALERT

### Description

This message occurs when WAFL® determines that there are duplicate Snapshot(tm) copy entries in the Snapshot copy translation table. Access to this Snapshot copy is disabled until the problem is resolved.

### Corrective Action

The Snapshot copy is inaccessible because the corresponding entry in the translation table has been cleared. Contact NetApp technical support resolve the problem.

## Syslog Message

WAFL Snapshot copy translation table: Duplicate entry for volume '%s%s%s' that maps physical Snapshot copy '%d' to logical Snapshot copy '%d'.

## Parameters

**owner** (STRING): Owner of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.  
**physical\_snapid** (INT): Physical Snapshot copy index.  
**Logical\_snapid** (INT): Logical Snapshot copy index.

# waf.snapshot\_xlation\_setupfailed events

## waf.snapshot\_xlation\_setupfailed

### Severity

ERROR

### Description

This message occurs when WAFL® cannot setup the Snapshot™ copy translation table. Access to this Snapshot copy is disabled until the problem is resolved.

### Corrective Action

The Snapshot copy is inaccessible, possibly due to a resource issue. Contact NetApp technical support.

## Syslog Message

WAFS Snapshot copy translation table: Cannot to add entry on volume '%s%s%s' that maps physical Snapshot copy '%d' to logical Snapshot copy '%d'

## Parameters

**owner** (STRING): Owner of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique Identifier of the volume in cases where the volume name itself is insufficient.

**physical\_snapid** (INT): Physical Snapshot copy index.

**Logical\_snapid** (INT): Logical Snapshot copy index.

# wafS.spacemgmt events

## wafS.spacemgmt.policyChg

### Severity

NOTICE

### Description

This message occurs when autodelete/autosize policy of a volume is changed.

### Corrective Action

(None).

## Syslog Message

The space management policy for volume %s has changed: %s.

## Parameters

**vol** (STRING): Volume name.

**changes** (STRING): A string of changes being made.

# wafS.spacetax events

## wafS.spacetax.disabled

### Severity

ERROR

### Description

This message occurs when the spacetax feature is disabled for all flexible volumes in the system because of the 'wafS-disable-spacetax?' boot option.

### Corrective Action

Reboot with the 'wafS-disable-spacetax?' boot option removed.

## Syslog Message

Spacetax disabled for all flexible volumes.

## Parameters

(None).

# wafl.sparse events

## wafl.sparse.meta.rm

### Severity

NOTICE

### Description

This event is issued when a sparse config metafile can not be removed in recycling bin

### Corrective Action

(None).

### Syslog Message

Unable to remove sparse metafile %s in recycle bin: error %s

### Parameters

**path** (STRING): The metadata filepath

**err** (STRING): The error message

# wafl.spcmgmt events

## wafl.spcmgmt.policyConflict

### Severity

NOTICE

### Description

This message occurs when a space management strategy or an autosize grow threshold for a volume is being configured and the autogrow or autodelete thresholds conflict with the space-mgmt-try-first policy. The conflict can arise when the autogrow threshold is less than the autodelete threshold and the space-mgmt-try-first policy is configured for "snap\_delete". The conflict can also arise when the autogrow threshold is more than the autodelete threshold and space-mgmt-try-first is configured for "volume\_grow".

### Corrective Action

Consider changing either the space-mgmt-try-first policy with the command "volume modify -volume <volume\_name> -space-mgmt-try-first <option>" or the autogrow threshold with the command "volume modify -volume <vol\_name> -autosize-grow-threshold-percent <value>", and keep them in sync.

### Syslog Message

While configuring "%s" for volume "%s", space\_mgmt\_try\_first policy ("%s") is conflicted with the autogrow threshold (%d%%) and autodelete threshold (%d%%).

### Parameters

**command** (STRING): Command being executed.

**volume\_name** (STRING): Name of the volume.

**space\_mgmt\_try\_first** (STRING): space-mgmt-try-first policy for volume\_grow or snap\_delete.

**grow\_threshold** (INT): Used-space threshold, as a percentage, for the automatic growth of the volume.  
**delete\_threshold** (INT): Used-space threshold, as a percentage, for the auto-deletion of Snapshot copies.

## wafl.spcrec events

### wafl.spcrec.not.changed

#### Severity

NOTICE

#### Description

This event indicates that the new space reservation guarantee was unable to be obtained, so the old reservation type was kept unchanged.

#### Corrective Action

(None).

#### Syslog Message

Retaining volume reservation on %s%s (%s) because new reservation could not be guaranteed.

#### Parameters

**vol\_name** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**spcrec\_desc** (STRING): Text string describing space reservation type.

## wafl.spinhi events

### wafl.spinhi.new.err.code

#### Severity

NOTICE

#### Description

This message occurs when a new spinnp error is being propagated to a blade that does not understand that error yet. Nothing bad should happen because the error will receive default error handling when received by that blade. This message is for debugging in case you need to know the original source and reason for the error.

#### Corrective Action

(None).

#### Syslog Message

Spinnp operation %s was executed on DSID(%08x), fileid %u, snapid %u, proto %s and returned spinnp error %d to blade UUID %s.'

#### Parameters

**opname** (STRING): Spinnp operation name.

**dsid** (LONGINT): The data set identifier.

**fileid** (INT): ID of the file specified by the spinnp operation.

**snapid** (INT): Snapshot(tm) copy identifier: (0) - active file system (non-zero) - refers to Snapshot copy

**proto** (STRING): Protocol used on the operation.

**spinerror** (INT): Spinnp error number.

**bladeuuid** (STRING): Universally Unique Identifier (UUID) of the blade where this error will be returned.

## wafl.split events

### wafl.split.infile.scan.complete

#### Severity

NOTICE

#### Description

This event indicates that the split infile scan has completed.

#### Corrective Action

(None).

#### Syslog Message

%s %s%s%s%s Split infile scan completed.

#### Parameters

**typestr\_cap** (STRING): wafl\_Volume type, capitalized

**owner** (STRING): owner string (foreign, partner, et al)

**vol\_name** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wafl.stream events

### wafl.stream.check.complete

#### Severity

NOTICE

#### Description

This message occurs when the "wafl scan stream\_check" command is completed successfully.

#### Corrective Action

(None).

#### Syslog Message

"wafl scan stream\_check" command successfully completed on volume "%s:%s". The total number of invalid stream files detected is %u.

#### Parameters

**volume** (STRING): Volume name.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's SVM.

**total** (INT): Total number of invalid stream files detected.

## **wafl.stream.check.invalid.file**

### **Severity**

ERROR

### **Description**

This message occurs when the "wafl scan stream\_check" command detects a stream file with the name "." or "..". Files with these names are not allowed. If you have files with these names, you might be unable to access these files.

### **Corrective Action**

Manually delete the base file which in turn will also delete the invalid stream file.

### **Syslog Message**

Detected invalid stream file associated with base file inode "%d" on volume "%s:%s".

### **Parameters**

**fileId** (INT): The base file inode number that has a stream file with the name "." or "..".

**volume** (STRING): Volume name.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's SVM.

## **wafl.stream.check.start**

### **Severity**

NOTICE

### **Description**

This message occurs when the "wafl scan stream\_check" command is started.

### **Corrective Action**

(None).

### **Syslog Message**

"wafl scan stream\_check" command successfully started on volume "%s:%s".

### **Parameters**

**volume** (STRING): Volume name.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the volume's SVM.

## **wafl.svo events**

### **wafl.svo.checkFailed**

#### **Severity**

NOTICE

#### **Description**

This event is generated if a write to a svo enabled volume fails to pass a check

## Corrective Action

(None).

## Syslog Message

### SnapValidator: Validation error %s

v:%s vol:%s%s%s inode:%d length:%d Offset: %lld

## Parameters

**error** (STRING): error that occurred when validating the write

**version** (STRING): string to describe version of oracle tests that were run

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**inode** (INT): The inode number of the file which was being written

**length** (INT): The length of the invalid write

**offset** (LONGINT): The offset into the inode where the write sent. Note for Luns this will be to the main inode, and does not include space used by prefix streams.

## waf1.svo.optionSet

### Severity

INFORMATIONAL

### Description

This event is generated when a SnapValidator option is set on a volume

### Corrective Action

(None).

## Syslog Message

SnapValidator option '%s' set to '%s' on volume '%s%s'

## Parameters

**opname** (STRING): name of the option being set

**opvalue** (STRING): value of the option being set

**vol** (STRING): name of the volume the option is being set on

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.takeover events

### waf1.takeover.nvram.missing

### Severity

INFORMATIONAL

### Description

This message occurs when the system does not find any WAFL® nvlog records to replay during takeover processing.



**Corrective Action**

(None).

**Syslog Message**

WAFL takeover: No WAFL nvlog records were found to replay.

**Parameters**

(None).

**waf1.takeover.nvram.warn****Severity**

ERROR

**Description**

This message occurs when NVRAM loss is detected during takeover. Some updates might have been lost, but takeover is proceeding.

**Corrective Action**

(Call support).

**Syslog Message**

WAFL takeover: Last few seconds of updates to partner might be lost.

**Parameters**

(None).

**waf1.takeover.vol.mType.inv****Severity**

ERROR

**Description**

This message occurs when WAFL® attempts to take over a mirror volume that does not have a valid mirror-type attribute. This situation might be a direct result of an upgrade to a version of Data ONTAP® that requires the mirror-type attribute or as a result of a corruption.

**Corrective Action**

This condition cannot be fixed by file system repair mechanisms such as waf1ron. The data is not lost, but it is inaccessible to the administrator. To access the data, contact NetApp technical support. Alternatively, the administrator can destroy the mirror volume by using the 'volume delete -vserver vserver\_name -volume vol\_name' command.

**Syslog Message**

The mirror type associated with the mirror volume %s%s%s is invalid.

**Parameters**

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name by itself is insufficient.

## **waf1.takeover.vol.nomount**

### **Severity**

ERROR

### **Description**

This message occurs during a takeover or switchover operation when the system discovers that a partner volume cannot be mounted because it is inconsistent.

### **Corrective Action**

The volume is inconsistent. Run waf1iron to clear this inconsistency, and then contact technical support.

### **Syslog Message**

%s partner volume %s%s%s is not mounted because it is inconsistent.

### **Parameters**

**operation** (STRING): Whether this message refers to takeover or switchover.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## **waf1.takeover.vol.offline**

### **Severity**

ERROR

### **Description**

This message occurs during a takeover or switchover operation when the system discovers that a partner volume cannot be brought online because it is inconsistent. The volume is left unmounted.

### **Corrective Action**

The volume is inconsistent and is left unmounted. Run waf1iron to clear this inconsistency, and then contact technical support.

### **Syslog Message**

%s partner volume %s%s%s cannot be brought online because it is inconsistent.

### **Parameters**

**operation** (STRING): Whether this message refers to takeover or switchover.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## **waf1.takeover.vols.mounted**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the volumes are mounted during takeover.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

(None).

## **waf1.timer.events**

### **waf1.timer.disk.io.stalled**

**Severity**

ERROR

**Description**

This message occurs when disk I/O is stalled, possibly during a disk firmware download operation. If automatic firmware update is enabled, the download operation restarts automatically.

**Corrective Action**

If disk firmware updates were in progress, verify that the update has automatically resumed by using the "storage download status" command. If the disk firmware updates have not resumed, restart the download by using the "storage download" command.

**Syslog Message**

All disk I/O stalled.

**Parameters**

(None).

## **waf1.trans.events**

### **waf1.trans.ssd.cache.disable**

**Severity**

NOTICE

**Description**

This message occurs when the system transitions a hybrid aggregate from 7-Mode to clustered Data ONTAP. As part of transition, SSD caching is disabled until the transition is committed.

**Corrective Action**

No action is required. After the transition operation is committed, SSD caching is automatically reenabled.

**Syslog Message**

SSD caching on hybrid aggregate "%s" has been disabled because a transition operation is in the pre-commit phase.

## Parameters

**name** (STRING): Aggregate name.

# wafl.transition events

## wafl.transition.cp.completed

### Severity

NOTICE

### Description

This message occurs at the completion of a WAFL® consistency point (CP) that occurs during a failover transition, such as takeover or giveback. This event logs the time spent and other important performance information.

### Corrective Action

(None).

### Syslog Message

Transition CP with reason %s for %s volumes with a total of %s dirty buffers took %ums with longest CP phases being %s on aggregate %s.

## Parameters

**tag** (STRING): Reason for the CP occurring and a status code.

**vols** (STRING): Counts of flexible volumes that participated in the CP.

**dirty\_bufs** (STRING): Number of WAFL buffers that were in use when the CP began and the number required to complete the CP.

**total\_time** (INT): Number of milliseconds that the CP ran for.

**top\_phases** (STRING): Top time consuming phases of the CP.

**aggrname** (STRING): Name of the aggregate.

# wafl.unable events

## wafl.unable.markinconsistent

### Severity

INFORMATIONAL

### Description

This message occurs when the FlexVol® volume cannot be marked as inconsistent.

### Corrective Action

(None).

### Syslog Message

WAFL: Unable to mark volume %s inconsistent as %s.

## Parameters

**vvolname** (STRING): Volume name.

**reason** (STRING): Reason for the inability to mark the volume inconsistent.

# wafl.upgrade events

## wafl.upgrade.useActiveInofile

### Severity

NOTICE

### Description

This event is issued during an upgrade where the system need to use the active filesystem inofile.

### Corrective Action

(None).

### Syslog Message

Using active filesystem inofile on upgrade from legacy %s %s%s%s%s.

### Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

# wafl.vol events

## wafl.vol.access.denied

### Severity

ERROR

### Description

This message occurs when the system rejects client access to a volume.

### Corrective Action

You can avoid this event by unmounting the volume.

### Syslog Message

Access to volume '%s' was denied, reason='%s'.

### Parameters

**vol** (STRING): Volume name.

**reason** (STRING): Reason that access was denied.

## wafl.vol.add

### Severity

INFORMATIONAL

## Description

This message indicates that a new volume has been added to the system.

## Corrective Action

(None).

## Syslog Message

%s %s%s%s%s has been added to the system.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## waf1.vol.add.keyinfo.into.vldb

### Severity

NOTICE

### Description

This message occurs when encryption key-info is added into vldb table.

### Corrective Action

(None).

### Syslog Message

Key-info is added into vldb for volume %s (dsid %lu, aggr '%s') with index %d, key-id %s.

### Parameters

**vol** (STRING): Volume name.

**dsid** (LONGINT): Unique identifier (DSID) of the volume.

**aggr** (STRING): Name of the containing aggregate.

**key\_idx** (INT): Key index.

**key\_id** (STRING): Key id.

## waf1.vol.appdm.upgrade.nospc

### Severity

ERROR

### Description

This message occurs when ONTAP® software cannot increase the size of an AppDM volume by five percent during an upgrade to ONTAP 9.5.

### Corrective Action

ONTAP has increased the Snapshot(tm) copy reserve percentage on this AppDM volume from zero to five percent as part of the upgrade to ONTAP 9.5. To retain the same available free space on the volume, consider increasing the size of the volume by five percent by using the "vol modify -size" command. It might be necessary to increase the size of the aggregate first, using the "aggr add-disks" command.

## Syslog Message

The file system on %s %s%s%s%s could not be grown 5%% larger as part of the upgrade to ONTAP 9.5.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier when the volume name alone is insufficient.

## waf1.vol.autoSize.done

### Severity

NOTICE

### Description

This message occurs when an autosize for a volume is successful. This can happen when the "autosize grow" option is on and the volume reaches the grow threshold percentage.

### Corrective Action

(None).

## Syslog Message

Volume autosize: Automatic %s of volume '%s%s%s' by %s is complete.

## Parameters

**event\_type** (STRING): Type of autosize event.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**size** (STRING): Human readable size.

**object\_uuid** (STRING): Volume instance UUID of the resource object.

## waf1.vol.autoSize.fail

### Severity

NOTICE

### Description

This message occurs when automatic sizing of a volume fails.

### Corrective Action

Increase the maximum size of the volume while setting the autosize.

## Syslog Message

Unable to grow volume '%s%s%s' to recover space: %s.

## Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier the volume in cases where the volume name itself is insufficient.

**err** (STRING): Reason for the failure.

**object\_uuid** (STRING): Volume instance UUID of the resource object.

## **waf1.vol.autoSize.shrink.cap**

### **Severity**

NOTICE

### **Description**

As a volume grows, there can be an automatic increase in the number of files that it is allowed to contain, but that number is unchanged as the volume shrinks. This number of files puts a constraint on the minimum size of the volume when it is automatically shrunk. This message indicates that the automatic shrinking of a volume did not reach the configuration's minimum size because it is constrained by the number of files the volume is allowed to contain.

### **Corrective Action**

Consider using the command "volume modify <volume-name> -files <new-files-value>" to decrease the number of files that the volume is allowed to contain. You can view the current value of the number of files by using the command "volume show <volume-name> -fields files".

### **Syslog Message**

Volume Autosize: Volume '%s%s%s' could not be auto shrunk below %s to recover space.

### **Parameters**

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**size** (STRING): Human-readable minimum size to which the volume can be automatically shrunk.

## **waf1.vol.bgTruncate.full**

### **Severity**

ALERT

### **Description**

This message occurs when a background truncation operation for large files stalls due to insufficient free space for allocating additional metadata on the volume.

### **Corrective Action**

Manually delete Snapshot(tm) copies or increase the volume's size for one-time space issues. For frequent space issues, consider changing autogrow or autodelete settings for the volume in question.

### **Syslog Message**

Insufficient free space available on %s %s%s %s to complete background truncation operation.

### **Parameters**

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier.



## **waf1.vol.blks\_used.done**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the 'blks used' calculation is complete.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**run\_time** (LONGINT): Total run time of the scanner.

## **waf1.vol.clearing.in.nvfailed.state**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when a volume that was marked as in-nvfailed-state is explicitly cleared.

### **Corrective Action**

(None).

### **Syslog Message**

Volume %s%s%s that was marked as in-nvfailed-state has been cleared.

### **Parameters**

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): If the volume belongs to a Vserver, the Vserver's Universal Unique Identifier (UUID).

## **waf1.vol.clone.split.nospace**

### **Severity**

NOTICE

### **Description**

This event is issued when there's not enough free space to complete a volume clone split operation. The usual cause for this is when there are many data blocks in the snapshot that is shared between the clone volume and its parent.

### Corrective Action

Free space in the clone volume to allow the split operation to take place, for example by deleting all snapshots in the clone volume.

### Syslog Message

Volume clone split of '%s%s' cannot be completed, not enough disk space. Consider removing snapshots.

### Parameters

**name** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf.vol.clone.unlock.snapshot.destroy.failed

### Severity

NOTICE

### Description

This event is issued when the volume options are set such that the autodelete destroy policy is allowed to destroy volume clones referring to the snapshot to be deleted. It indicates an attempt to destroy a volume clone has failed.

### Corrective Action

(None).

### Syslog Message

Volume clone destroy of '%s' failed.

### Parameters

**name** (STRING): The volume clone name

## waf.vol.clone.unlock.snapshot.failure

### Severity

NOTICE

### Description

This event is issued when the volume options are set such that the autodelete destroy policy is allowed to destroy volume clones referring to the snapshot to be deleted. It indicates the process of destroying volume clones has failed. Some volume clones may have been destroyed or left offline as a result.

### Corrective Action

(None).

### Syslog Message

Process of destroying volume clones using volume %s snapid %d failed. Some volume clones may have been taken offline or destroyed.

### Parameters

**name** (STRING): The volume name

**snapid** (INT): Snapshot(tm) copy identifier.

## **waf1.vol.clone.unlock.snapshot.offline.failed**

### **Severity**

NOTICE

### **Description**

This event is issued when the volume options are set such that the autodelete destroy policy is allowed to destroy volume clones referring to the snapshot to be deleted. It indicates an attempt to offline a volume clone has failed.

### **Corrective Action**

(None).

### **Syslog Message**

Volume clone offline of '%s' failed.

### **Parameters**

**name** (STRING): The volume clone name

## **waf1.vol.clone.unlock.snapshot.start**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the volume options are set so that the autodelete destroy policy is allowed to destroy volume clones referring to the Snapshot™ copy to be deleted. The message indicates that the process of destroying volume clones is beginning. The Snapshot autodelete destroy policy is used when all other attempts to recover space in a thin-provisioned aggregate fail. The disruptive effect of destroying volume clones should be understood when setting this policy.

### **Corrective Action**

(None).

### **Syslog Message**

Starting process of destroying volume clones using volume %s snapid %d.

### **Parameters**

**name** (STRING): Volume name.

**snapid** (INT): Snapshot copy identifier.

## **waf1.vol.clone.unlock.snapshot.success**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the volume options are set so that the autodelete destroy policy is allowed to destroy volume clones referring to the Snapshot™ copy to be deleted. The message indicates that the process of destroying volume clones has successfully completed.

**Corrective Action**

(None).

**Syslog Message**

Destroyed all volume clones using using volume %s snapid %d.

**Parameters**

**name** (STRING): Volume name.

**snapid** (INT): Snapshot copy identifier.

**waf.vol.create.clusterVolOnCFOaggr****Severity**

ERROR

**Description**

This message occurs when a volume is created on a Controller Failover (CFO) high-availability (HA) style aggregate. File access delays will occur for a period following giveback.

**Corrective Action**

To minimize file access delays following giveback, create the volume on an aggregate with a Storage Failover (SFO) HA policy.

**Syslog Message**

Creating volume '%s' on an aggregate with a CFO HA policy. File access delays will occur for a period following giveback.

**Parameters**

**vol** (STRING): Name of the volume.

**waf.vol.deswizzle.done****Severity**

INFORMATIONAL

**Description**

This message occurs when the deswizzle scan is complete.

**Corrective Action**

(None).

**Syslog Message**

(None).

**Parameters**

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**run\_time** (LONGINT): Total run time of the scanner.

## waf1.vol.fix.comp.algo.list

### Severity

NOTICE

### Description

This message occurs when CP detects and rectifies an incorrect compression algorithm list in volume.

### Corrective Action

(None).

### Syslog Message

List of used compression algorithms in the volume %s has been fixed. was %llu, now %llu.

### Parameters

**vol** (STRING): Volume name.

**old\_val** (LONGINT): Old value of compression algorithm list

**new\_val** (LONGINT): New value of compression algorithm list

## waf1.vol.fsp.full

### Severity

ERROR

### Description

This message occurs when an operation requests more space than is available in the Flexible Space Pool.

### Corrective Action

Report this condition to NetApp technical support.

### Syslog Message

%s %s%s%s%s: insufficient space in FSP %s to satisfy a request of %llu holes and %llu overwrites.

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier, in order to uniquely identify the volume in cases where volume name alone is insufficient.

**fspname** (STRING): The name of the flexible space pool.

**requested** (LONGINT): Amount of space requested from the flexible space pool.

**requested\_overwrites** (LONGINT): Amount of overwrite reserve space requested from the flexible space pool.

## waf1.vol.full

### Severity

ALERT

## Description

This message occurs when an operation requests more space than is available in the relevant volume.

## Corrective Action

Consider manually deleting data or increasing the volume's size for one-time space issues. For frequent space issues, consider changing autogrow/autodelete settings.

## Syslog Message

Insufficient space on %s %s%s%s%s to perform operation. %s was requested but only %s was available.

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Volume identifier, in order to uniquely identify the volume in cases where volume name alone is insufficient.

**requested** (STRING): Amount of space requested in the volume.

**available** (STRING): Amount of space available in the volume.

## waf1.vol.guarantee.fail

### Severity

NOTICE

### Description

This message occurs when an attempt is made to apply the space guarantee for a flexible volume, but there is not enough space in the aggregate to do so. The volume will be mounted, but subsequent writes to the volume might fail if the volume's aggregate becomes full.

### Corrective Action

Either increase the size of the corresponding aggregate by adding disks, or decrease the total volume space requirements in the aggregate by shrinking or destroying one or more volumes.

### Syslog Message

Space for volume %s%s is not guaranteed. %s of additional space is required to enable the guarantee.

### Parameters

**vol** (STRING): Name of the volume whose space guarantee could not be applied.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**short\_by** (STRING): Amount of additional space required in the volume.

## waf1.vol.guarantee.stayoff

### Severity

NOTICE

### Description

This message occurs when WAFL® would otherwise have attempted to enable a volume's space guarantee, but the last time it tried, an error was encountered (most likely due to insufficient space). This error was recorded in the volume's on-disk RAID label so that the same error would not be repeated on

future mounts. The volume will be mounted. However, because the volume is not guaranteed, it is possible that subsequent writes to the volume could fail if the volume's aggregate becomes full.

### Corrective Action

Determine the cause of the error the last time WAFL attempted to enable this guarantee, either through logs or by attempting to enable it again manually. Typically, the issue is a lack of space at mount time detected during boot, which can be resolved by increasing the size of the corresponding aggregate (by adding disks or by shrinking or destroying some of the aggregate's volumes). After the issue is resolved, manually attempt to re-enable the guarantee using the 'volume modify' command.

### Syslog Message

The space guarantee for volume %s%s was not enabled because an error occurred the last time a guarantee was enabled on this volume.

### Parameters

**vol** (STRING): Name of the volume.  
**app** (STRING): Application UUID.  
**volident** (STRING): Volume identifier.

## waf.vol.inode.space.usage

### Severity

ERROR

### Description

This message occurs when public inode space usage exceeds threshold value configured by admin.

### Corrective Action

Use the command "volume modify <volume-name> -files <new-files-value>" to increase the number of files that the volume is allowed to contain. View the current value of the file threshold by using the command "volume show <volume-name> -fields files". If volume modify fails, increase volume size and try again.

### Syslog Message

The file system on %s %s%s%s%s is currently using %d%% of the files that can be contained on the volume. The admin has configured the threshold for monitoring public inode usage as %d%%.

### Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identifier when the volume name alone is insufficient.  
**current\_infile\_used** (INT): Current value of inodes in public inode space.  
**threshold\_ems\_infile\_used\_pct** (INT): Threshold percent value to monitor inodes usage.

## waf.vol.invalid.nomount

### Severity

NOTICE

## Description

This event is issued when a volume has not been fully copied and cannot be mounted.

## Corrective Action

(None).

## Syslog Message

Volume %s is not fully copied and cannot be mounted.

## Parameters

**vol** (STRING): The volume name

## waf1.vol.invalidate.fh

## Severity

NOTICE

## Description

This message indicates that an NVRAM failure is detected, invalidating outstanding filehandles.

## Corrective Action

(None).

## Syslog Message

%s %s%s%s%s: Invalidating all filehandles due to previous NVRAM failure — Filesystem must be remounted by the client(s)

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## waf1.vol.mirrorType.invalid

## Severity

ERROR

## Description

This message occurs when WAFL® attempts to bring online a mirror volume that does not have a valid mirror-type attribute. This situation might be a direct result of an upgrade to a version of Data ONTAP® that requires the mirror-type attribute or as a result of a corruption.

## Corrective Action

This condition cannot be fixed by file system repair mechanisms such as waf1ron. The data is not lost, but it is inaccessible to the administrator. To access the data, contact NetApp technical support. Alternatively, the administrator can destroy the mirror volume by using the 'volume delete -vserver vservice\_name -volume vol\_name' command.



## Syslog Message

The mirror type associated with the mirror volume %s %s%s%s%s is invalid.

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name by itself is insufficient.

## waf.vol.mislabeled.root

### Severity

NOTICE

### Description

This event is issued when a volume or aggregate is labeled as root but was not selected as the root volume or aggregate. It is used to indicate that the label will be removed.

### Corrective Action

(None).

## Syslog Message

%s '%s%s%s%' was incorrectly labeled as root. Removing the label.

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner Can be the empty string

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf.vol.nvfail.offline

### Severity

EMERGENCY

### Description

This message occurs when the system detects an NVRAM failure and the volume is inconsistent. The volume might have data loss due to NVRAM failure. The volume has taken offline.

### Corrective Action

Take the appropriate recovery action to recover from the possible data loss. Access to the volume is restricted until you clear the NVFAIL state manually.

## Syslog Message

Volume %s%s is offline.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

## waf1.vol.nvfail.recovery.needed

### Severity

NOTICE

### Description

This message occurs periodically to inform the administrator that recovery procedures must be performed on a volume that is marked as in-nvfailed-state.

### Corrective Action

Take the appropriate recovery action to recover from the possible data loss. Access to the volume is restricted until you clear the NVFAIL state manually.

### Syslog Message

Volume %s%s%s is marked as in-nvfailed-state. The recovery process needs to be initiated on the volume to recover from the NVFAIL event.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): If the volume belongs to a Vserver, the Vserver's Universal Unique Identifier (UUID).

## waf1.vol.offline

### Severity

INFORMATIONAL

### Description

This message indicates that a volume has been taken offline.

### Corrective Action

(None).

### Syslog Message

%s %s%s%s%s has been taken temporarily offline.

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**instuuid** (STRING): Volume Instance UUID

## waf1.vol.offline.callbacks

### Severity

NOTICE

### Description

Data ONTAP was trying to offline a volume and had to ignore the result of file-locking callbacks being sent to clients in order to make progress. This is transparent to clients, as they will retry their original requests,

which will be serviced when the volume comes back online.

### Corrective Action

(None).

### Syslog Message

Ignoring %d callbacks when offlining %s %s%s%s%s.

### Parameters

**count** (INT): The number of callbacks which were ignored

**type** (STRING): The type of object (volume or aggregate)

**owner** (STRING): The owner of the volume

**vol** (STRING): The name of the volume

**app** (STRING): Application UUID.

**volident** (STRING): The unique identifier of the volume in case the volume name itself is insufficient

## waf.vol.outOfInodes

### Severity

ERROR

### Description

This message occurs when a volume or aggregate reaches its maximum number of files and is out of inodes.

### Corrective Action

For a volume, either use the command "volume modify -files" to increase the number the files that the volume is allowed to contain, or delete files from the volume. For an aggregate, add space to the aggregate.

### Syslog Message

file system on %s %s%s%s%s is out of inodes because it has reached the maximum number of files

### Parameters

**type** (STRING): Type of object (volume or aggregate)

**owner** (STRING): Owner of the volume or aggregate

**vol** (STRING): Name of the volume or aggregate

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume or aggregate in cases where the name itself is insufficient.

## waf.vol.outOfPrInodes

### Deprecated

### Severity

NOTICE

### Description

This event is issued when the system discover that a volume is out of private inodes.

### Corrective Action

(None).

## Syslog Message

file system on %s %s%s%s%s is out of private inodes

## Parameters

**type** (STRING): The type of object (volume/aggregate)

**owner** (STRING): The volume owner

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf.vol.quiesce.giveback

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a volume quiescing has failed because waf giveback is running.

### Corrective Action

Retry quiescing the volume after giveback has completed.

## Syslog Message

Quiescing volume %s failed because waf giveback is running.

## Parameters

**vol** (STRING): The volume name

## waf.vol.quiesce.iron

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a volume quiescing has failed because waf iron is running.

### Corrective Action

Retry quiescing the volume after waf iron has completed.

## Syslog Message

Quiescing volume %s failed because waf iron is running.

## Parameters

**vol** (STRING): The volume name

## waf.vol.quiesce.notdrain

### Severity

INFORMATIONAL

## Description

This message occurs when a volume quiesce operation fails because of a general failure in draining the suspended messages or outstanding callbacks. This covers unknown cases that `waf.vol.quiesce.timeoutcb` does not catch.

## Corrective Action

(None).

## Syslog Message

Quiescing of volume %s failed because draining message and callbacks failed. The total time spent quiescing the volume was %d milliseconds, %d of which was spent aborting scanners, %d milliseconds were spent flushing client requests for locks and watches, %d milliseconds were spent draining volume references and %d milliseconds were spent draining callbacks and suspended messages.

## Parameters

**vol** (STRING): Volume name.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.

**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.

**twdcm** (INT): Time spent (in milliseconds) waiting for draining of suspended messages and callbacks.

## waf.vol.quiesce.notFlex

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a volume quiescing has failed because it is not a flex vol.

### Corrective Action

(None).

### Syslog Message

Quiescing volume %s failed because it's not a flexible volume.

### Parameters

**vol** (STRING): The volume name

## waf.vol.quiesce.notMounted

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a volume quiescing has failed because it is not in mounted state.

### Corrective Action

(None).

## Syslog Message

Quiescing volume %s failed because it's not mounted. It's in %s state.

## Parameters

**vol** (STRING): The volume name

**statename** (STRING): The state

## waf.vol.quiesce.relocating

### Severity

ERROR

### Description

This message occurs when a volume cannot be quiesced because an aggregate relocation operation is running.

### Corrective Action

Retry the volume move operation after the aggregate relocation operation is complete.

## Syslog Message

Volume %s could not be quiesced because aggregate relocation is running.

## Parameters

**vol** (STRING): Volume name.

## waf.vol.quiesce.rootVol

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a volume quiescing has failed because it is a root volume.

### Corrective Action

(None).

## Syslog Message

Quiescing volume %s failed because it's a root volume.

## Parameters

**vol** (STRING): The volume name

## waf.vol.quiesce.statusvol

### Severity

ERROR

### Description

This message occurs when a volume cannot be quiesced because it is being used to host wafiron status files.

### Corrective Action

Use "aggr wafiron show" command to monitor the status of wafiron and retry the operation after wafiron is complete.

### Syslog Message

Volume %s could not be quiesced because it is being used by wafiron to host status files.

### Parameters

**vol** (STRING): Volume name.

## waf.vol.quiesce.timeoutcb

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring draining of suspended callbacks or messages fails because of a timeout.

### Corrective Action

(None).

### Syslog Message

Quiescing of volume %s failed because there are %d outstanding callbacks and %d suspended messages. The time spent waiting for these to drain was %d milliseconds. The time spent quiescing the volume (%d milliseconds) was split between scan aborts (%d), flushing client requests (%d), and draining references (%d).

### Parameters

**vol** (STRING): Volume name.

**ncbo** (INT): Number of outstanding callbacks.

**nsm** (INT): Number of suspended messages.

**twdcm** (INT): Time spent (in milliseconds) waiting for draining of suspended messages and callbacks.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.

**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.

## waf.Vol.QuiesceTO.AbrtScan

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring scanners to abort fails because of a timeout.

### Corrective Action

(None).

## Syslog Message

Quiescing of the volume %s failed because aborting the scanners resulted in a timeout. The time spent quiescing the volume was %d milliseconds of which %d milliseconds were spent waiting for the scanners to abort.

## Parameters

**vol** (STRING): Volume name.

**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

## waf1.Vol.QuiesceTO.CIntReqs

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring flushing of client requests for locks and watches fails because of a timeout.

### Corrective Action

(None).

## Syslog Message

Quiescing of volume %s failed because time spent waiting for flushing of client requests for locks and watches (%d milliseconds) resulted in a timeout. The total time spent quiescing the volume was %d milliseconds - %d milliseconds were spent aborting scanners.

## Parameters

**vol** (STRING): Volume name.

**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.

## waf1.Vol.QuiesceTO.RalDrain

### Severity

NOTICE

### Description

This message occurs during a volume move operation, when a volume quiesce attempt fails to drain all of the outstanding remote cached inodes within the allowable time frame. This can occur and might persist on a heavily loaded system, and it prevents the move from being completed. However, as soon as the load drops to a more reasonable level, ONTAP flushes the cached inodes and completes the move without customer intervention.

### Corrective Action

(None).

## Syslog Message

Quiescing of volume %s failed because it was unable to drain all outstanding remote inodes for the volume. Initially the volume reported %d cached inodes, but timed out with %d cached inodes remaining.



## Parameters

**vol** (STRING): Volume name.  
**nStart** (INT): Number of cached inodes when drain starts.  
**nEnd** (INT): Number of cached inodes when request times out.

## waf1.Vol.QuiesceTO.RefDrain

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring draining of volume references fails because of a timeout.

### Corrective Action

(None).

### Syslog Message

Quiescing of volume %s failed because the time spent draining the references on the volume (%d milliseconds) resulted in a timeout. There are %d references on the volume. One of the owners is %s. The command "waf1 vref status" under diag privileges can be used to view all the references on a volume on the node owning the volume. The time spent quiescing the volume (%d milliseconds) was split between scan aborts (%d) and flushing client requests(%d).

## Parameters

**vol** (STRING): Volume name.  
**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.  
**wvc** (INT): Number of references on the volume.  
**vo\_1** (STRING): Text string representing the first volume reference owner.  
**ttq** (INT): Time spent (in milliseconds) quiescing the volume.  
**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.  
**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

## waf1.Vol.QuiesceTO.RefsDrain

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring draining of volume references fails because of a timeout. The references are released.

### Corrective Action

(None).

### Syslog Message

Quiescing of volume %s failed because the time spent draining the references on the volume (%d milliseconds) resulted in a timeout. The time spent quiescing the volume (%d milliseconds), was split between scan aborts (%d) and flushing client requests (%d).

## Parameters

**vol** (STRING): Volume name.

**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

**twc** (INT): Time spent (in milliseconds) waiting for scanners to abort.

**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

## waf1.Vol.QuiesceTO.SAN.Ntfy

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring a notification to the SAN component fails because of a timeout.

### Corrective Action

(None).

### Syslog Message

Quiescing of volume %s failed because a quiesce notification to the SAN component (%d milliseconds) resulted in a timeout. The time spent quiescing the volume (%d milliseconds) was split between scan aborts (%d), flushing client requests (%d), draining references (%d), draining callbacks and suspended messages (%d), and draining zombies (%d).

## Parameters

**vol** (STRING): Volume name.

**twsc** (INT): Time spent (in milliseconds) waiting to for quiescing of the SAN component.

**ttq** (INT): Time spent (in milliseconds) quiescing the volume.

**twc** (INT): Time spent (in milliseconds) waiting for scanners to abort.

**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.

**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.

**twdcm** (INT): Time spent (in milliseconds) waiting for draining of suspended messages and callbacks.

**twzd** (INT): Time spent (in milliseconds) waiting for draining of zombies.

## waf1.Vol.QuiesceTO.ZmbDrain

### Severity

INFORMATIONAL

### Description

This message occurs when a volume quiesce operation requiring draining of zombies fails because of a timeout.

### Corrective Action

(None).

### Syslog Message

Quiescing of volume %s failed because the time spent draining zombies (%d milliseconds) resulted in a timeout. The time spent quiescing the volume (%d milliseconds) was split between scan aborts (%d), flushing client requests (%d), draining references (%d), and draining callbacks and suspended messages (%d).

## Parameters

**vol** (STRING): Volume name.  
**twzd** (INT): Time spent (in milliseconds) waiting for draining of zombies.  
**ttq** (INT): Time spent (in milliseconds) quiescing the volume.  
**tws** (INT): Time spent (in milliseconds) waiting for scanners to abort.  
**twcr** (INT): Time spent (in milliseconds) flushing client requests for locks and watches.  
**twvr** (INT): Time spent (in milliseconds) waiting for draining of volume references.  
**twdcm** (INT): Time spent (in milliseconds) waiting for draining of suspended messages and callbacks.

## waf.vol.rekey.stat

### Severity

NOTICE

### Description

This message occurs when rekey scan ends to dump the stats.

### Corrective Action

(None).

### Syslog Message

Rekey scan is complete on volume '%s' (dsid %lu, aggr '%s'). Scan Stats: max\_usecs\_in\_cpu = %lu, blks\_total = %lu, max\_blks\_rekeyed\_in\_step = %lu, blks\_rekeyed\_by\_scanner = %lu, blks\_rekeyed\_by\_client\_io = %lu, blks\_with\_error = %lu, blks\_skipped\_dirty = %lu, blks\_suspended\_fake\_dirty = %lu, blks\_skipped\_holes = %lu, blks\_packed = %lu

## Parameters

**vol** (STRING): Volume name.  
**dsid** (LONGINT): Unique identifier (DSID) of the volume.  
**aggr** (STRING): Name of the containing aggregate.  
**max\_usecs\_in\_cpu** (LONGINT): Max. time (in micro-sec.) scanner runs in CPU.  
**blks\_total** (LONGINT): Total number of blocks scanner is expected to process.  
**max\_blks\_rekeyed\_in\_step** (LONGINT): Max. number of blocks scanner process in one step.  
**blks\_rekeyed\_by\_scanner** (LONGINT): Number of blocks rekeyed by scanner.  
**blks\_rekeyed\_by\_client\_io** (LONGINT): Number of blocks rekeyed by client I/O.  
**blks\_with\_error** (LONGINT): Number of blocks failed to be loaded by scanner due to error.  
**blks\_skipped\_dirty** (LONGINT): Number of blocks skipped by scanner because they are already dirty.  
**blks\_suspended\_fake\_dirty** (LONGINT): Number of blocks rekey scanner suspended on because they were fake dirty.  
**blks\_skipped\_holes** (LONGINT): Number of blocks skipped by scanner because they are holes.  
**blks\_packed** (LONGINT): Number of blocks packed.

## waf.vol.runningOutOfNodes

### Severity

ERROR

### Description

This message occurs when the system detects that a volume is close to reaching its limit on the number of files that it can contain.

## Corrective Action

Consider using the "volume modify <volume-name> -files <new-files-value>" command to increase the number of files that the volume is allowed to contain. You can view the current value of the file threshold by using the "volume show <volume-name> -fields files" command.

## Syslog Message

The file system on "%s %s%s%s%s" is using 80 percent or more of the files in %s inode space that can be contained on the volume.

## Parameters

**ino\_type** (STRING): Inode type

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique volume identifier when the volume name alone is insufficient.

## waf.vol.setting.in.nvfailed.state

### Severity

ERROR

### Description

This message occurs when the system detects an NVFAIL event in a volume and marks the volume as in-nvfailed-state.

### Corrective Action

Access to the volume is restricted until you clear the NVFAIL state manually. See the Data Protection Online Backup and Recovery Guide.

### Syslog Message

Volume %s%s%s is marked as in-nvfailed-state: %s.

### Parameters

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): If the volume belongs to a Vserver, the Vserver's Universal Unique Identifier (UUID).

**err** (STRING): Error that caused the volume to be marked in in-nvfailed-state.

## waf.vol.snap\_create.done

### Severity

INFORMATIONAL

### Description

This message occurs when the 'snap create' scan is complete.

### Corrective Action

(None).

## Syslog Message

(None).

## Parameters

**type** (STRING): Type of object (volume or aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**run\_time** (LONGINT): Total run time of the scanner.

## waf1.vol.spcsl0.conflict

### Severity

ERROR

### Description

This message occurs when the Service Level Objective related to space management (the space SLO setting) for a volume conflicts with other volume settings.

### Corrective Action

Try correcting the conflicting volume setting or space SLO by using the "volume modify <volname>" command.

## Syslog Message

The space SLO setting %s for volume %s%s is incompatible with the volume's %s setting (%s)%s.

## Parameters

**space\_slo** (STRING): Space SLO of the volume.

**vol** (STRING): Name of the volume with conflicting settings.

**app** (STRING): Application UUID.

**volident** (STRING): String that uniquely identifies the volume when the volume name is insufficient.

**setting** (STRING): Name of the conflicting volume setting.

**value** (STRING): Value of the conflicting volume setting.

**detail** (STRING): Corrective action already taken, if any.

## waf1.vol.tooFullForUnlink

### Severity

ALERT

### Description

This message occurs when an unlink operation requires more space than is available for unlinking in the relevant volume. Clients cannot remove files or directories until free space is added.

### Corrective Action

Manually delete Snapshot(tm) copies or increase the volume's size for one-time space issues. For frequent space issues, consider changing autogrow or autodelete settings for the volume or aggregate in question.

## Syslog Message

Insufficient free space available on %s %s %s%s %s to perform unlink operation.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Volume identifier.

## waf1.vol.unable.autogrow

### Severity

ERROR

### Description

This message occurs when automatically growing a volume fails because the underlying storage of the SVM (Vserver) has reached the SVM's storage limit.

### Corrective Action

Increase the storage of the SVM by raising the storage limit. Alternatively, recover storage space by deleting volumes, Snapshot copies, LUNs, or files.

### Syslog Message

Auto-grow of volume "%s" failed because the SVM's storage limit has been reached.

## Parameters

**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.  
**object\_uuid** (STRING): Volume instance UUID of the resource object.

## waf1.vol.walloc.rsv.dbg.reset

### Severity

NOTICE

### Description

This event is issued to indicate that the write allocation reserve thresholds for a volume or aggregate have reset to normal. This message is issued on DEBUG builds only.

### Corrective Action

(None).

### Syslog Message

DBG: Volume %s allocation reserve thresholds reset.

## Parameters

**vol** (STRING): The volume or aggregate name

## waf1.vol.walloc.rsv.dbg.warn.blkcnt

**Severity**

NOTICE

**Description**

This event is issued to indicate that the write allocation reserve code has triggered but the volume block counters appear to be in the normal range.

**Corrective Action**

(None).

**Syslog Message**

DBG: Volume (or aggregate) %s free space reserve hit %s threshold but block count doesn't match.

**Parameters**

**vol** (STRING): The volume or aggregate name

**type** (STRING): Threshold type: "low" or "high"

**waf1.vol.walloc.rsv.failmount****Severity**

EMERGENCY

**Description**

This message occurs when a volume consumes almost all of its free space reserve and should not be brought back online without assistance from technical support to address the free space issue.

**Corrective Action**

Contact NetApp technical support.

**Syslog Message**

Failing mount of volume %s. Free space reserve is too low. Contact NetApp technical support for assistance.

**Parameters**

**vol** (STRING): Volume name.

**waf1.vol.walloc.rsv.offline****Severity**

ALERT

**Description**

This message occurs when a volume or aggregate has consumed almost all of its free space reserve. The system is being shut down and rebooted cleanly and the affected aggregate or volume will be taken offline until the situation can be corrected with help from technical support.

**Corrective Action**

Contact NetApp technical support..

**Syslog Message**

Volume %s free space reserve is critically low. Rebooting system to take volume offline. Contact technical

support for assistance.

### Parameters

**vol** (STRING): Volume or aggregate name.

## waf.vol.walloc.rsv.root

### Severity

ALERT

### Description

This message occurs when the root volume has consumed almost its entire free space reserve. This is a highly abnormal condition that can lead to data loss if the free space reserve is completely consumed.

### Corrective Action

Increasing the root volume size. Contact NetApp technical support for more information.

### Syslog Message

Root volume (or aggregate) %s free space reserve is critically low.

### Parameters

**vol** (STRING): Root volume or aggregate name.

## waf.volinfo events

### waf.volinfo.fsinfo.error

### Severity

ALERT

### Description

This message occurs when there is an error while reading the volinfo or fsinfo block of a volume.

### Corrective Action

Take appropriate action based on the reason for the error.

### Syslog Message

%s for %s '%s'

### Parameters

**reason** (STRING): Reason for the error

**type** (STRING): The type of object (volume/aggregate)

**vol** (STRING): The name of the volume

## waf.volmove events

### waf.volmove.dst.amd.corrupt



**Severity**

ALERT

**Description**

This message occurs when the active map delta file for a volume move destination is corrupt.

**Corrective Action**

Run wafiron on the volume after the "volume move" operation is complete. For further assistance, contact NetApp technical support.

**Syslog Message**

Active map delta file for the "volume move" destination "%s" is corrupt.

**Parameters**

**vol** (STRING): Volume name.

## waf.volref events

### waf.volref.subsystem.string

**Severity**

ERROR

**Description**

This message occurs when wafiron is started on aggregate containing volumes, which are in-use by other subsystems.

**Corrective Action**

Stop the subsystem which is using this volume/aggregate So that wafiron can be run.

**Syslog Message**

wafiron: %s %s is being used by '%s'

**Parameters**

**type** (STRING): Type of object (volume or aggregate).

**vol** (STRING): Name of the volume or aggregate.

**volref** (STRING): volume reference string of the subsystem using this volume/aggregate.

## waf.voltrans events

### waf.voltrans.clear.noi2p

**Severity**

NOTICE

**Description**

This message occurs when a volume transition clears the noi2p option on a volume.

**Corrective Action**

(None).

**Syslog Message**

Volume transition will disable the noi2p option by default on volume %s.

**Parameters**

**volname** (STRING): Name of the volume.

**waf.voltrans.clear.svo****Severity**

NOTICE

**Description**

This message occurs when a volume transition operation clears the SnapValidator for Oracle options on a volume.

**Corrective Action**

(None).

**Syslog Message**

Volume transition will disable the SnapValidator® for Oracle (svo) options by default on volume %s.

**Parameters**

**volname** (STRING): Name of the volume.

**waf.voltrans.dircopy.fail****Severity**

ALERT

**Description**

This message occurs when a volume transition convert directory operation fails.

**Corrective Action**

Check the status of the job by using the "volume transition-convert-dir show" (privilege: diag) command. Retry the volume transition after all volume directory transition operations have finished. If the operation fails on the second attempt, contact NetApp technical support.

**Syslog Message**

Volume transition-convert-dir operation failed on volume "%s", copying path "%s" to temporary directory "%s". Error code: %d.

**Parameters**

**volname** (STRING): Name of the volume.

**path** (STRING): Path of the directory to be copied.

**tempname** (STRING): Directory name to use as the temporary target for copy conversion.

**errcode** (INT): Error code upon completion.

## **waf1.voltrans.flexcache**

### **Severity**

ERROR

### **Description**

This message occurs during volume transition workflow when the transition volume is a FlexCache® volume. Support for FlexCache is deprecated.

### **Corrective Action**

Delete the FlexCache volume.

### **Syslog Message**

Volume transition failed due to unsupported FlexCache volume %s uuid %s.

### **Parameters**

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## **waf1.voltrans.inconsistent.state**

### **Severity**

ERROR

### **Description**

This message occurs when a volume transition operation fails because the volume is in an inconsistent state.

### **Corrective Action**

Delete the inconsistent volume. Alternatively, return the disks to a 7-Mode environment to repair the volume.

### **Syslog Message**

Volume transition failed for volume %s uuid %s due to its inconsistent state.

### **Parameters**

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## **waf1.voltrans.qtree.novfiler**

### **Severity**

NOTICE

### **Description**

This message occurs when a volume transition operation encounters qtrees not associated with the default vFiler unit.

### **Corrective Action**

(None).

## Syslog Message

Volume transition detected non-default vfiler qtrees in volume %s%s%s.

## Parameters

**volname** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the owning Vserver when the volume name alone is insufficient.

## waf1.voltrans.qtree.replica

### Severity

ERROR

### Description

This message occurs when a volume transition operation is attempted on a volume that is a qtree replica volume. Transitioning volumes that are qtree replica volumes is an unsupported operation. This condition exists because the source volume for the Transition Data Protection (TDP) SnapMirror® relationship is a qtree replica volume.

### Corrective Action

Break all of the qtree replica relationships associated with the source volume of the TDP SnapMirror relationship. After that is done, perform a 'snapmirror update' or 'snapmirror resync' operation, and then a 'snapmirror break' operation, to attempt the volume transition operation again.

## Syslog Message

The source of volume %s is a qtree replica. It cannot be transitioned to clustered Data ONTAP.

## Parameters

**volname** (STRING): Name of the volume.

## waf1.voltrans.sfsr.pending

### Severity

ERROR

### Description

This message occurs when a volume transition operation fails because of a pending file restore process.

### Corrective Action

Move the disks back to 7-Mode, use ZAPI command 'snapshot-restore-file-info' to ensure that all pending single file restore sessions are completed, and then try the transition again.

## Syslog Message

Volume transition failed for volume %s uuid %s due to pending single file restore process.

## Parameters

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## **waf.voltrans.snapids**

### **Severity**

ERROR

### **Description**

This message occurs when a volume transition operation fails because it cannot update the logical SnapID table.

### **Corrective Action**

(None).

### **Syslog Message**

Volume %s failed to update the logical SnapID table and could not be transitioned.

### **Parameters**

**volname** (STRING): Name of the volume.

## **waf.voltrans.snapshot.32bit**

### **Severity**

ERROR

### **Description**

This message occurs during volume transition workflow when 32-bit Snapshot copy is detected for the transition volume.

### **Corrective Action**

Retry volume transition to bypass this specific error. All 32-bit Snapshot copies are automatically deleted as part of the volume transition.

### **Syslog Message**

Volume transition failed due to unsupported 32-bit snapshots on volume %s uuid %s.

### **Parameters**

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## **waf.voltrans.ss.lun.clone**

### **Severity**

ERROR

### **Description**

This message occurs when a volume transition operation detects one or more LUN clones backed up by Snapshot® copies within a volume.

### **Corrective Action**

Retry volume transition to bypass this specific error, force the volume online, and continue the transition operation. This results in data loss for all the LUN clones that have Snapshot copies in this volume. Alternatively, return the disks to a 7-Mode environment, split the LUN clones, and then repeat the transition workflow.

## Syslog Message

Volume transition failed for volume %s uuid %s due to one or more snapshot backed lun clones.

## Parameters

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## waf.voltrans.vol.32bit

### Severity

ERROR

### Description

This message occurs during volume transition workflow when a 32-bit volume is detected for the transition volume.

### Corrective Action

Delete the 32-bit volume. Alternatively, return the disks to a 7-Mode environment and upgrade the 32-bit volume to a 64-bit volume and then repeat the transition workflow.

## Syslog Message

Volume transition failed due to unsupported 32-bit volume %s uuid %s.

## Parameters

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## waf.voltrans.was.offline

### Severity

ERROR

### Description

This message occurs when a volume transition operation encounters an offline volume.

### Corrective Action

Retry volume transition to bypass this specific error, force the volume online, and continue the transition operation. This might lead to data loss if the offline volume contains uncommitted data.

## Syslog Message

Volume transition failed for offline volume %s uuid %s.

## Parameters

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

## waf.voltrans.was.restricted

### Severity

ERROR

## Description

This message occurs when a volume transition operation encounters a restricted volume.

## Corrective Action

Retry volume transition to bypass this specific error, force the volume online, and continue the transition operation.

## Syslog Message

Volume transition failed for restricted volume %s uuid %s.

## Parameters

**volname** (STRING): Name of the volume.

**voluuid** (STRING): Unique identifier for this volume.

# wafv.volume events

## wafv.volume.charmap.loadfail

### Severity

ERROR

### Description

This message occurs when a volume comes online with a corrupted or unavailable character mapping configuration.

### Corrective Action

Remove character mapping by using the "vserver cifs character-mapping delete" command. Then reapply it by using the "vserver cifs character-mapping create" command. For more information, see the "vserver cifs character-mapping" man page. For further assistance, contact NetApp technical support.

### Syslog Message

Failed to load character mapping configuration for volume %s. Reason: %s.

### Parameters

**volname** (STRING): Volume for which loading of the character mapping configuration failed.

**reason** (STRING): Reason for failure.

## wafv.volume.clone.created

### Severity

INFORMATIONAL

### Description

This message occurs when a clone of a flexible volume is created successfully.

### Corrective Action

(None).

### Syslog Message

Volume clone %s of volume %s was created successfully.

## Parameters

**clonename** (STRING): Volume clone name

**volname** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## wafl.volume.clone.fractional\_rsrv.changed

### Severity

INFORMATIONAL

### Description

This message occurs when the space reservation of a flexible volume clone is changed to 100 percent because guarantee of the clone is set to 'file'.

### Corrective Action

(None).

### Syslog Message

Fractional reservation for clone '%s:%d' was changed to 100 percent because guarantee is set to 'file'.

## Parameters

**clonename** (STRING): Volume clone name.

**Ordinal** (INT): Volume Ordinal.

## wafl.volume.clone.snap.cpcount.load.failure

### Severity

ERROR

### Description

This message occurs when a clone of a flexible volume fails to load the cp count of its or the parent volume's base snapshots.

### Corrective Action

(None).

### Syslog Message

Volume clone %s failed to load the cp count of its or the parent volume's base snapshot, with Snap ID %d.

## Parameters

**clonename** (STRING): Volume clone name.

**snapid** (INT): Snapshot ID of the clone volume's base snapshot.

## wafl.volume.clone.split.started

### Severity

INFORMATIONAL



## Description

This message occurs when an operation to split a flexible volume clone from its parent volume is started successfully.

## Corrective Action

(None).

## Syslog Message

Clone split was started for volume %s%s%s

## Parameters

**volname** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## wafl.volume.clone.with.uid.gid

### Severity

INFORMATIONAL

### Description

This message occurs when a clone of a flexible volume is created with a volume-level UID and GID.

### Corrective Action

Split the clone volume before creating this as a source of an XDP-type SnapMirror relationship.

### Syslog Message

Volume clone %s%s%s was created with vol-level UID and GID. If this clone volume is used as a source of XDP type SnapMirror relationship, then UID and GID of all the transferred file at the destination will be unpredictable.

### Parameters

**clonename** (STRING): Volume clone name.

**app** (STRING): Application UUID.

**vserver\_uuid** (STRING): Universal Unique ID (UUID) of the clone volume's Vserver.

## wafl.volume.find.valid\_base.snap.failure

### Severity

NOTICE

### Description

This event is issued when an error occurs while finding the valid base snapshot for a clone volume.

### Corrective Action

None required.

### Syslog Message

Encountered error %u while finding the valid base snapshot having snapid %u for clone volume (clonename='%s', volident='%s').

## Parameters

**error\_code** (INT): A numerical error code indicating the reason why finding valid base snapshot failed for clone volume.

**snapid** (INT): Snapshot ID of the clone volume base snapshot.

**clonename** (STRING): Volume clone name.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## wafv.volume.inuse

### Severity

ALERT

### Description

This message occurs when a volume operation fails because the volume is being used in another operation.

### Corrective Action

Check the status of the volume and aggregate using the "vol show" and "aggr show" commands. If there are no other operations running on either of them then retry your request. If problems persists, or you need further assistance, contact NetApp technical support.

### Syslog Message

%s %s %s operation failed as volume is busy in another operation.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**vol** (STRING): Name of the volume.

**operation** (STRING): Operation.

## wafv.volume.inuse.by

### Severity

ALERT

### Description

This message occurs when a volume operation fails because the volume or parent aggregate is being used by another operation.

### Corrective Action

Check the status of the volume and aggregate. Wait for the current operation to finish, and then retry your request. If problems persist or you need further assistance, contact NetApp technical support.

### Syslog Message

%s %s %s operation failed because the volume or parent aggregate is locked by operation %s.

### Parameters

**type** (STRING): Type of object (volume or aggregate).

**vol** (STRING): Name of the volume.

**operation** (STRING): Operation that failed.

**locked\_by** (STRING): Operation currently locking the volume or parent aggregate.

## **waf1.volume.sidl.enableErr**

### **Severity**

ERROR

### **Description**

This message occurs when an attempt to enable Single Instance Data Logging (SIDL) on a volume fails.

### **Corrective Action**

Use the "volume efficiency show" and "volume show" commands to check the volume for enabled features that are incompatible with SIDL, such as compaction or inline deduplication. After the incompatible features are disabled, wait 10 seconds and then enable SIDL again.

### **Syslog Message**

SIDL was not enabled on volume %s

### **Parameters**

**volume** (STRING): Volume name.

## **waf1.volume.sidl.offline**

### **Severity**

ERROR

### **Description**

This message occurs when a volume with single instance data logging enabled is taken offline due to the current version of Data ONTAP not supporting that feature. This situation can occur if a volume with the feature enabled is attempted to be onlined using a version of Data ONTAP that doesn't support it.

### **Corrective Action**

To bring the volume online, use a version of Data ONTAP that supports single instance data logging.

### **Syslog Message**

Volume "%s" was taken offline because the current version of Data ONTAP does not support single instance data logging, an enabled feature on this volume.

### **Parameters**

**volume** (STRING): Volume name.

## **waf1.volume.slag.loadfailed**

### **Severity**

ERROR

### **Description**

This message occurs when a loading or update of Storage-Level Access Guard (SLAG) fails because of an internal error. In this case, SLAG is not activated on the volume or qtree.

### **Corrective Action**

Remove SLAG by using the "vserver security file-directory remove-slag" command, and then reapply it by using the "vserver security file-directory apply" command.

## Syslog Message

Failed to load or update SLAG for SLAG level: %d on qtree ID: %d in volume "%s%s%s%s". Error:%d  
Reason: %s.

## Parameters

**loadflag** (INT): SLAG load flag.

**qtree** (INT): Qtree identifier.

**owner** (STRING): Owner of the volume.

**vol** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume in cases where the volume name itself is insufficient.

**waferror** (INT): WAFL error which caused the failure.

**reason** (STRING): Reason for the failure.

## wafl.volume.snap.autoDelete

### Severity

INFORMATIONAL

### Description

This message occurs when a Snapshot(tm) copy in the volume is being deleted to free space. Snapshot copies can be deleted automatically from a volume to enforce the corresponding automatic deletion policy.

### Corrective Action

(None).

## Syslog Message

Deleting Snapshot copy '%s' in volume '%s%s%s' to recover storage.

## Parameters

**name** (STRING): Snapshot copy name.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## wafl.volume\_snap events

### wafl.volume\_snap.autoDelete.fail

### Severity

INFORMATIONAL

### Description

This message occurs when automatic deletion of a Snapshot(tm) copy fails.

### Corrective Action

(None).

## Syslog Message

Unable to delete Snapshot copy '%s' in volume '%s%s%s' to recover space: %s

## Parameters

**name** (STRING): Snapshot copy name.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Uniquely identifies the volume when the volume name itself is insufficient.

**err** (STRING): Reason for failure.

# wafv.vv events

## wafv.vv.bad.DSIDorMSID

### Severity

ERROR

### Description

This message occurs when Data ONTAP® cannot map a volume to serve a client request because of invalid data Set Identifiers. Protocol access to the volume is disabled.

### Corrective Action

Take the volume offline and contact NetApp technical support.

### Syslog Message

Cannot map volume '%s%s%s' to serve client requests because of invalid data set identifiers.

### Parameters

**vol** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wafv.vv.clone.splitBusySnap

### Severity

NOTICE

### Description

This event is to inform the user that a volume clone cannot split from its parent volume due to busy snapshots.

### Corrective Action

Determine which snapshots in the clone volume are busy and correct the condition that is causing them to be busy. Then, retry the operation.

### Syslog Message

Volume '%s%s%s' has locked snapshots, not splitting.

### Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.vv.clone.splitNoScan

### Severity

NOTICE

### Description

This event is to inform user that the system is unable to start a volume clone scanner, so the clone volume will not be split from its parent volume.

### Corrective Action

This may be a temporary issue. Please retry the volume clone split operation after sometime.

### Syslog Message

Unable to start clone split scanner on volume '%s%s%s'.

### Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.vv.dup.DSID2

### Severity

ERROR

### Description

This message occurs when another volume is already registered with the data set identifier (DSID) assigned to this volume.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Duplicate volume '%s%s%s' DSID (%08x) detected in aggregate '%s'.

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**dsid** (LONGINT): Data set identifier.

**aggr** (STRING): Aggregate name.

## waf1.vv.dup.UUID

### Severity

NOTICE

### Description

This message occurs when the system detects and updates a duplicate volume Universally Unique Identifier (UUID).

**Corrective Action**

(None).

**Syslog Message**

vvol '%s' in aggregate '%s' duplicate UUID %s was changed to %s.

**Parameters**

**vol** (STRING): Volume name.

**aggr** (STRING): Aggregate name.

**old\_uuid** (STRING): Old duplicated UUID for the volume.

**new\_uuid** (STRING): Replacement UUID for the volume.

**wafv.vv.raidfile.open****Severity**

NOTICE

**Description**

This event is to inform user that the system can not open the raid label file for update and the reason.

**Corrective Action**

(None).

**Syslog Message**

Unable to open volume descriptor: %s

**Parameters**

**err** (STRING): The error message

**wafv.vv.raidfile.readOnly****Severity**

NOTICE

**Description**

This event is to inform user that the raid label file can not be updated on a read only volume.

**Corrective Action**

(None).

**Syslog Message**

Unable to update volume descriptor, aggregate '%s' is readonly

**Parameters**

**vol** (STRING): The volume name

**wafv.vv.raidfile.update****Severity**

NOTICE

## Description

This event is to inform user that while renaming a virtual volume, the system is unable to update the volume's raid label file

## Corrective Action

(None).

## Syslog Message

%s: Can't rewrite descriptor for volume '%s%s%s', aggregate '%s', error code %d.

## Parameters

**cmd** (STRING): The command name

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**aggr** (STRING): The aggregate name

**err** (INT): The error code

## wafv.vv.regen.instance.uuid

### Severity

INFORMATIONAL

### Description

This message occurs when the instance UUID and the provenance UUID of a FlexVol(tm) volume have been regenerated.

### Corrective Action

(None).

### Syslog Message

For FlexVol volume '%s' with FSID 0x%x and instance UUID %s, a new instance UUID and provenance UUID (both being %s) have been generated.

### Parameters

**vol** (STRING): Volume name.

**fsid** (INT): File system identifier.

**oldinstuuid** (STRING): Old instance UUID of the volume.

**instuuid** (STRING): New instance UUID of the volume. This is also the new provenance UUID.

## wafv.vv.rename.dup

### Severity

NOTICE

### Description

This message occurs when a duplicate flexible volume name is changed to a unique name.

### Corrective Action

(None).



## Syslog Message

Duplicate flexible volume name '%s%s%s' changed to '%s%s%s' in aggregate '%s'.

## Parameters

**oldname** (STRING): Duplicate name.

**app** (STRING): Application UUID.

**volident1** (STRING): If the volume belongs to a vserver, this is that vserver's Universal Unique Identifier (UUID).

**newname** (STRING): New name.

**app2** (STRING): Application UUID.

**volident2** (STRING): If the volume belongs to a vserver, this is that vserver's Universal Unique Identifier (UUID).

**aggrname** (STRING): Name of the containing aggregate.

## waf1.vv.set.size

### Severity

NOTICE

### Description

This event is issued when setting virtual volume size to the size indicated in fsinfo

### Corrective Action

(None).

## Syslog Message

Setting volume size for %s%s%s%s to match nominal size.

## Parameters

**owner** (STRING): The volume owner

**name** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## waf1.vv.softlocks.noParent

### Severity

NOTICE

### Description

This event is to inform user that to remove softlocks associated with a volume clone, the parent volume must be online

### Corrective Action

(None).

## Syslog Message

Unable to find parent volume when destroying clone '%s%s%s' softlocks. Ensure the parent volume is online then retry the operation.

## Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

## wafv.vol.config

### Severity

NOTICE

### Description

This event is issued to inform the user that the system can not configure a virtual volume and the reason.

### Corrective Action

(None).

### Syslog Message

Unable to configure volume in aggregate '%s': %s

## Parameters

**vol** (STRING): The volume name

**err** (STRING): The error string

## wafv.vol.maxvols

### Severity

NOTICE

### Description

This message occurs when the system cannot configure a FlexVol® volume because the maximum number of volumes has been reached.

### Corrective Action

Destroy some existing volumes to make room for new ones.

### Syslog Message

Cannot configure FlexVol volume '%s' (UUID %s) in aggregate '%s': There are too many FlexVol volumes (%d) on the system. The limit for %s is %d.

## Parameters

**vol** (STRING): FlexVol name.

**uuid** (STRING): FlexVol unique ID.

**aggregate** (STRING): Aggregate name.

**count** (INT): Current number of FlexVol volumes.

**model** (STRING): System model name.

**limit** (INT): Limit for the number of FlexVol volumes.

## wafv.vvlock events

## waf1.vvclock.string

### Severity

ERROR

### Description

This message occurs when an operation is run on an aggregate while the flexible volume level admin (vvc) lock of one of volumes of the aggregate is held by another subsystem. The message prints the name the subsystem that is holding the lock and the type of lock.

### Corrective Action

Stop the subsystem that is using this volume/aggregate so that the operation can be run.

### Syslog Message

%s: aggregate %s is busy. Reason: vvc lock of the volume %s is held by subsystem %s. %s.

### Parameters

**operation** (STRING): Name of the operation.

**aggrname** (STRING): Name of the aggregate, one of the volumes of which is holding the vvc lock.

**volname** (STRING): Name of the volume that is holding the vvc lock.

**lockerstring** (STRING): Name of the subsystem that is holding the vvc lock.

**locktype** (STRING): vvc lock type:WAFL® or SnapMirror®.

## waf1.vvol events

### waf1.vvol.create.encrypted

### Severity

NOTICE

### Description

This message occurs when an encrypted flexible volume is created.

### Corrective Action

(None).

### Syslog Message

New %s encrypted volume %s is created on aggregate %s.

### Parameters

**key\_type** (STRING): This Indicates encryption type of this volume is NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE).

**vol** (STRING): Volume name.

**aggr** (STRING): Name of the containing aggregate.

**voluuid** (STRING): Volume UUID.

**keyid** (STRING): Encryption Key ID.

**bufreeid** (INT): Bufreeid.

**cp\_count** (INT): Consistency Point (CP) at which this volume is created.

## **waf1.vvol.delete.metadata.recrm**

### **Severity**

NOTICE

### **Description**

This message occurs when the system cannot remove volume metadata during a volume delete operation after moving it to recycle bin.

### **Corrective Action**

(None).

### **Syslog Message**

Could not remove volume %s%s%s metadata. Error: %d.

### **Parameters**

**vol** (STRING): Name of the flexible volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name alone is insufficient.

**errorCode** (INT): Error number.

## **waf1.vvol.delete.metadata.rename**

### **Severity**

NOTICE

### **Description**

This message occurs when the system cannot rename volume metadata during a volume delete operation.

### **Corrective Action**

(None).

### **Syslog Message**

Could not rename volume %s%s%s metadata. Error: %d.

### **Parameters**

**vol** (STRING): Name of the flexible volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name alone is insufficient.

**errorCode** (INT): Error number.

## **waf1.vvol.delete.metadata.rm**

### **Severity**

NOTICE

### **Description**

This message occurs when the system cannot remove volume metadata during a volume delete operation. The system will retry removing the metadata during any operation that brings the containing aggregate online(for example, takeover).

## Corrective Action

(None).

## Syslog Message

Cannot remove volume %s%s%s metadata. Error: %s(%d).

## Parameters

**vol** (STRING): Name of the flexible volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name alone is insufficient.

**reason** (STRING): Reason for the failure.

**errorCode** (INT): Error number.

## waf.vvol.delete.recovered

### Severity

INFORMATIONAL

### Description

This event occurs when a flexible volume that was queued for delete is recovered.

### Corrective Action

(None).

### Syslog Message

Volume %s%s%s%s recovered.

### Parameters

**owner** (STRING): String indicating the owner of the affected volume.

**vol** (STRING): Name of the flexible volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name alone is insufficient.

## waf.vvol.destroyed

### Severity

NOTICE

### Description

This message occurs when a flexible volume is destroyed.

### Corrective Action

(None).

### Syslog Message

Volume %s%s%s%s destroyed.

### Parameters

**owner** (STRING): String indicating the owner of the affected volume.

**vol** (STRING): Name of the flexible volume.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier for the volume when the volume name alone is insufficient.  
**vol\_bt tid** (INT): Internal buftree identifier for the volume.

## **waf l.vvol.encrypted.offline**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when an encrypted virtual volume is taken offline.

### **Corrective Action**

(None).

### **Syslog Message**

Volume '%s%s%s' is encrypted and was taken offline.

### **Parameters**

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

## **waf l.vvol.excessive.contsize**

### **Severity**

ERROR

### **Description**

This message occurs when a flexible volume container that exceeds the maximum size supported on the running platform is detected. This event can be encountered after a headswap to a platform supporting a lower maximum capacity.

### **Corrective Action**

Back up the dataset in the volume as soon as possible and take the volume offline. Alternatively, consider a system or hardware upgrade to a platform that supports a higher maximum capacity.

### **Syslog Message**

Volume '%s' exceeded the maximum container capacity supported on this platform (container blkcnt=%llu, maxlimit=%llu). Back up the dataset and take the volume offline as soon as possible.

### **Parameters**

**vol** (STRING): The volume name

**blkcount** (LONGINT): The total blocks capacity of the volume container

**maxblkcount** (LONGINT): The maximum container capacity supported on the platform

## **waf l.vvol.excessive.volsize**

### **Severity**

ERROR

## Description

This message occurs when a flexible volume that exceeds the maximum size supported on the running platform is detected. This event can be encountered after a headswap to a platform supporting a lower maximum capacity.

## Corrective Action

If possible, shrink the flexible volume using the 'vol size' command. If not, back up the dataset in the volume as soon as possible and take the volume offline. Alternatively, consider a system or hardware upgrade to a platform that supports a higher maximum capacity.

## Syslog Message

Volume '%s' exceeded the maximum volume size supported on this platform (nominal blkcnt=%llu, maxlimit=%llu). Shrink the volume or back up the dataset and take the volume offline as soon as possible.

## Parameters

**vol** (STRING): Volume name

**blkcount** (LONGINT): Total block count of the volume.

**maxblkcount** (LONGINT): Maximum block count supported on the platform.

## waf1.vvol.invalid.DSID

### Severity

ERROR

### Description

This message occurs when ONTAP® finds a volume with an invalid data set identifier (DSID). This might have occurred as a result of disk corruption.

### Corrective Action

Contact NetApp technical support to have a new DSID assigned to the volume as soon as possible.

### Syslog Message

Volume '%s%s%s' has an invalid DSID (0x%x).

### Parameters

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**dsid** (LONGINT): Data set identifier.

## waf1.vvol.move.encrypted

### Severity

NOTICE

### Description

This message occurs when an encrypted destination volume is created as part of volume move operation.

### Corrective Action

(None).

## Syslog Message

%s encrypted volume %s%s%s is created on aggregate %s as part of a volume move operation.

## Parameters

**key\_type** (STRING): This Indicates encryption type of this volume is NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE).

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

**aggr** (STRING): Name of the containing aggregate.

**voluuid** (STRING): Volume UUID.

**keyid** (STRING): Encryption Key ID.

**buftreeid** (INT): Buftreeid.

**cp\_count** (INT): Consistency Point (CP) at which this volume is created.

## waf1.vvol.offline

### Severity

INFORMATIONAL

### Description

This message indicates when a virtual volume is taken offline.

### Corrective Action

(None).

## Syslog Message

Volume '%s%s%s' has been set temporarily offline

## Parameters

**name** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**instuuid** (STRING): Volume Instance UUID

## waf1.vvol.offline.err

### Severity

NOTICE

### Description

This event is issued if the system fail to set the raid label of the virtual volume when it's taken offline.

### Corrective Action

(None).

## Syslog Message

Error persistently transitioning volume %s%s%s from 'restricted' to 'offline', unable to update descriptor, error %d



## Parameters

**name** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**error** (INT): The error code

## waf1.vvol.offline.persist

### Severity

INFORMATIONAL

### Description

This message occurs when a flexible volume that was marked temporarily offline as part of the volume migration process is now marked persistently offline.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**name** (STRING): Volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient.

## waf1.vvol.online

### Severity

INFORMATIONAL

### Description

This event occurs when the state of flexible volume moves to the online state.

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vol** (STRING): The name of the flexible volume

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**instuuid** (STRING): Volume instance UUID

## waf1.vvol.ownVserver.changed

### Severity

INFORMATIONAL

## Description

This message occurs when either a cluster repair tool changes the owning vserver of the volume, or when the volume is transitioned from being a node volume to a cluster volume or vice versa.

## Corrective Action

(None).

## Syslog Message

Volume '%s%s' owning vserver / name changed to '%s%s%s'.

## Parameters

**oldname** (STRING): Original name of the volume.  
**oldownerIdent** (STRING): Original vserver identifier.  
**newname** (STRING): New name of the volume.  
**app** (STRING): Application UUID.  
**newownerIdent** (STRING): New vserver identifier.

## waf1.vvol.renamed

### Severity

INFORMATIONAL

### Description

This message occurs when the administrator renames a flexible volume.

### Corrective Action

(None).

### Syslog Message

Volume '%s%s' renamed to '%s%s%s'.

### Parameters

**oldname** (STRING): Original volume name.  
**volident1** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.  
**newname** (STRING): New volume name.  
**app** (STRING): Application UUID.  
**volident2** (STRING): Unique identifier of the volume in cases where the volume name itself is insufficient.

## waf1.vvol.restrict

### Severity

INFORMATIONAL

### Description

This event is issued to indicate a flexible volume has been restricted

### Corrective Action

(None).

### Syslog Message

(None).

## Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**instuuid** (STRING): Volume Instance UUID

## waf1.vvol.vvlabel.error

### Severity

ERROR

### Description

This message occurs when the verification of the vvlabel fails on a flexible volume. This event provides all the information available in the vvlabel so that it can be used later, for recovery.

### Corrective Action

Contact NetApp technical support.

### Syslog Message

Cannot determine volume label information (%s). (Flexible volume with UUID '%s', unique ID '%d' on aggregate '%s' with raidtree ID '%s')

## Parameters

**err** (STRING): Error string.

**voluuid** (STRING): UUID of the flexible volume as identified by the metadir.

**id\_in\_aggr** (INTEX): Unique ID within the containing aggregate.

**aggr\_name** (STRING): Containing aggregate name.

**aggr\_rtid** (STRING): RAID tree ID of the aggregate.

**version** (INTEX): vvlabel version.

**mirror\_flags** (INTEX): Mirror type flags.

**vol\_flags** (INTEX): Flags for the volume.

**fsid** (INTEX): File system identifier for the volume.

**clone\_parent** (INTEX): 'id\_in\_aggr' value of the clone's parent volume.

**base\_snapid** (INTEX): Clone base Snapshot(tm) copy ID.

**vvuuid** (STRING): UUID of the flexible volume as identified by the vvlabel.

**vvol\_type** (INTEX): Flexible volume's type.

**dsid** (LONGINTEX): Volume data set identifier.

**msid** (LONGINTEX): Volume master data set identifier.

**owneruuid** (STRING): UUID of the owning Vserver.

**volname** (STRINGARRAY): Volume name in raw hexadecimal format.

## waf1.vvol.xfer.nospacedestErr

### Severity

ERROR

### Description

This event indicates that the transfer to a destination flexible volume has failed because the aggregate than contains this volume is out of space.

## Corrective Action

Increase space available in the aggregate by either adding disks or freeing space by deleting aggregate snapshots, shrinking flexible volumes or moving flexible volumes to another aggregate. Ensure that there is sufficient space available in the aggregate (using the `df -A` command) to accommodate the flexible volumes it contains. The `vol size` command may be used to obtain flexible volume sizes and also to grow or shrink them.

## Syslog Message

Transfer to volume '%s%s%s' failed because the containing aggregate, '%s', is out of space.

## Parameters

**vol** (STRING): The volume name

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where volume name itself is insufficient.

**aggr** (STRING): The aggregate name

# wafL.watch events

## wafL.watch.cache.limit.near

### Severity

ERROR

### Description

This message occurs when WAFL® "change notify" requests approach the size limit of the inode cache pool. This is caused by an excessive number of "change notify" requests being generated by the Server Message Block(SMB) server. If the maximum limit is reached and the limit is not increased, significant system performance degradation could occur.

### Corrective Action

Contact Contact NetApp technical support. for assistance.

### Syslog Message

Size of "change notify" requests in inode cache pool: %d, approaching the size limit: %d

### Parameters

**current\_size** (INT): Current size of the inode cache pool.

**max\_size** (INT): Maximum size of the inode cache pool.

# wafL.write events

## wafL.write.fail.cifsSpace

### Severity

ERROR

### Description

This message occurs when a write to a file with CIFS-style space reservations fails due to a lack of space in the volume or aggregate.

## Corrective Action

If the volume is a Flexvol® volume, and the volume guarantee is disabled, increase the amount of space available in the aggregate. Otherwise, increase the amount of space available in the volume.

## Syslog Message

Write failed to file with CIFS space reservations due to lack of disk space in %s %s%s%s%s (%s%s, inode %u, offset %llu, len %u).

## Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**space\_info** (STRING): Text string specifying whether there is a guarantee.

**guarantee\_state** (STRING): Text string describing the state of the guarantee on the volume (disabled, volume, file, and so on).

**ino** (INT): Inode number.

**offset** (LONGINT): Block offset.

**len** (INT): Length value.

## waf1.write.fail.spcres

### Severity

ERROR

### Description

This message occurs when a write to a space-reserved file fails due to a lack of space in the volume or aggregate. This is possible when the fractional reserve is not 100% or the space guarantee of the containing flexible volume is not guaranteed.

### Corrective Action

If the volume is a Flexvol® volume, and the volume guarantee is disabled, increase the amount of space available in the aggregate. Otherwise, increase the amount of space available in the volume.

### Syslog Message

Write failed to file with space reservations due to lack of disk space in %s %s%s%s%s (%s%s, inode %u, offset %llu, len %u).

### Parameters

**type** (STRING): Type of object (volume/aggregate).

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

**volident** (STRING): Unique identifier of the volume when the volume name by itself is insufficient.

**space\_info** (STRING): Text string specifying whether there is a guarantee.

**guarantee\_state** (STRING): Text string describing the state of the guarantee on the volume (disabled, volume, file, and so on).

**ino** (INT): Inode number.

**offset** (LONGINT): Block offset.

**len** (INT): Length value.

# waf1.zombie events

## waf1.zombie.susp.aggr.limit

### Severity

NOTICE

### Description

This message occurs when the aggregate encounters the zombie throttling limits during zombie delete processing for large number of file deletions.

### Corrective Action

(None).

### Syslog Message

WAF1® aggregate %s is experiencing zombie throttling possibly due to requests for large number of file deletions. This can be mitigated by a combination of a) reducing the load on the aggregate or b) issuing the file deletion requests in smaller batches.

### Parameters

**aggrname** (STRING): Name of the aggregate.

## waf1.zombie.susp.vol.limit

### Severity

NOTICE

### Description

This message occurs when the volume encounters the zombie throttling limits during zombie delete processing for large number of file deletions.

### Corrective Action

(None).

### Syslog Message

WAF1® volume %s%s%s%s is experiencing zombie throttling possibly due to requests for large number of file deletions. This can be mitigated by a combination of a) reducing the load on the volume or b) issuing the file deletion requests in smaller batches.

### Parameters

**volowner** (STRING): Owner of the volume.

**volname** (STRING): Name of the volume.

**app** (STRING): Application UUID.

**volident** (STRING): To uniquely identify the volume in cases where the volume name itself is insufficient. Consists of the Vserver name and UUID.

# wafiron events

## wafiron.summary events

### wafiron.summary.scanStatus

#### Severity

INFORMATIONAL

#### Description

This message occurs periodically when wafiron is running and its state is "starting" or "scanning". It reports wafiron summary scanning status.

#### Corrective Action

(None).

#### Syslog Message

Wafiron summary scan (%d%% done).

#### Parameters

**percentage** (INT): Wafiron summary scan completion percentage.

# xfcvi events

## xfcvi.attr events

### xfcvi.attr.changed

#### Severity

NOTICE

#### Description

This message occurs when any of the xfcvi attributes changes.

#### Corrective Action

(None).

#### Syslog Message

FC-VI adapter: %s changed on port %s. %s

#### Parameters

**attr\_type** (STRING): Attribute type.

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

## xfcvi.cm events

### xfcvi.cm.timedout

#### Severity

NOTICE

#### Description

This message occurs when connection requests times out.

#### Corrective Action

(None).

#### Syslog Message

FC-VI adapter: Connection request timedout on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

#### Parameters

**port** (STRING): xfcvi Interconnect NIC.

**qpname** (STRING): QP name.

**qpnum** (INT): QP index.

**system\_id** (LONGINT): Remote node's system id.

**info** (STRING): Additional info.



## xfcvi.cross events

### xfcvi.cross.connect.detected

#### Severity

ERROR

#### Description

This message occurs when the xfcvi driver detects a cross connection

#### Corrective Action

Check the xfcvi interconnect cabling at the local and remote nodes. xfcvi adapter ports xfcvi\_device\_0s and xfcvi\_device\_1s should not be connected to the same fabric. Contact NetApp technical support if the problem persists.

#### Syslog Message

FC-VI adapter: Cross Connection detected between local port %s and remote port %s. Local System ID = %llu, Remote System ID = %llu.

#### Parameters

**lname** (STRING): Local node's port name.

**rname** (STRING): Remote node's port name.

**lsid** (LONGINT): Local node's system id.

**rsid** (LONGINT): Remote node's system id.

## xfcvi.disabled events

### xfcvi.disabled

#### Severity

ALERT

#### Description

This message occurs when the HA interconnect card is disabled by the ic-type environment variable.

#### Corrective Action

Either remove the HA interconnect card from the controller or change the ic-type environment variable.

#### Syslog Message

FC-VI adapter: Port %s has been disabled because of the ic-type environment variable.

#### Parameters

**port** (STRING): xfcvi Interconnect NIC.

## xfcvi.fcrenabled events

### xfcvi.fcrEnabled

**Severity**

NOTICE

**Description**

This message occurs if Fiber Channel Routing (FCR) functionality is enabled on the system.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Fiber Channel Routing (FCR) is enabled.

**Parameters**

(None).

## **xfcvi.initfailed events**

### **xfcvi.initFailed**

**Severity**

ERROR

**Description**

This message occurs when the driver fails to initialize a FC-VI port.

**Corrective Action**

Reboot the controller. Contact NetApp technical support if the problem persists.

**Syslog Message**

FC-VI adapter: %s failed for port %s.

**Parameters**

**operation** (STRING): Type of operation  
**port** (STRING): xfcvi Interconnect NIC.

## **xfcvi.invalid events**

### **xfcvi.invalid.advertisement**

**Severity**

ERROR

**Description**

This message occurs when xfcvi driver receives invalid advertisement.

**Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Port %s received invalid advertisement. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

## xfcvi.invalid.sfp

### Severity

ERROR

### Description

This message occurs when the SFP (Small Form-Factor Pluggable) module of the Fibre Channel Virtual Interface (FC-VI) is invalid or incompatible.

### Corrective Action

Replace the SFP with a compatible one.

## Syslog Message

FC-VI adapter: Invalid SFP - Reason: %s on port %s.

## Parameters

**reason** (STRING): Reason for the incompatibility.

**port** (STRING): Interconnect NIC.

## xfcvi.link events

### xfcvi.link.down

### Severity

NOTICE

### Description

This message occurs when the physical link of the Fibre Channel Virtual Interface (FC-VI) adapter goes down.

### Corrective Action

1. Ensure that the FC-VI link has not been tampered with.
2. Verify that the physical status of the FC-VI adapter is "Up" by using the command "metrocluster interconnect adapter show".
3. If the configuration includes fabric switches, ensure that they are properly cabled and configured.

## Syslog Message

(None).

## Parameters

**alertId** (STRING): Alert ID.

**alertingResource** (STRING): Alerting resource unique name.

**alertingResourceName** (STRING): Alerting resource name.

**possibleEffect** (STRING): Possible effect for link transition.

**probableCause** (STRING): Probable cause of the event.

**probableCauseDescription** (STRING): Probable cause description of the event.  
**additionalInfo** (STRING): Any additional information for the event.

## **xfcvi.link.up**

### **Severity**

NOTICE

### **Description**

This message occurs when the physical link of the Fibre Channel Virtual Interface (FC-VI) adapter comes up.

### **Corrective Action**

(None).

### **Syslog Message**

(None).

### **Parameters**

**alertId** (STRING): Alert ID.

**alertingResource** (STRING): Alerting resource unique name.

## **xfcvi.qlgc events**

### **xfcvi.qlgc.async.event.overflow**

#### **Severity**

ERROR

#### **Description**

This message occurs when total number of async events are overflowed.

#### **Corrective Action**

(None).

#### **Syslog Message**

FC-VI adapter: Port %s async event overflowed. Resetting the adapter.

#### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.

### **xfcvi.qlgc.connErr**

#### **Severity**

ERROR

#### **Description**

This message occurs when the FC-VI adapter detects a connection error. The FC-VI driver will attempt to retry the connection.

## Corrective Action

If this message repeats frequently, there may be an issue with the FC fabric or the adapter. Contact NetApp technical support.

## Syslog Message

FC-VI adapter: Port %s detects connection error. FC-VI opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**opcode\_str** (STRING): xfcvi opcode.  
**opcode** (INHEX): xfcvi opcode.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.cpr.reset.cancelled**

### Severity

ERROR

### Description

This message occurs when the CPR reset initiated by the FC-VI port got cancelled.

### Corrective Action

Reboot the filer. If issue still persists then replace the xfcvi adapter.

### Syslog Message

FC-VI adapter: CPR reset initiated by port %s got cancelled. Offlining the port.

### Parameters

**port** (STRING): xfcvi Interconnect adapter.

## **xfcvi.qlgc.cpr.reset.failed**

### Severity

ERROR

### Description

This message occurs when the CPR reset initiated by the FC-VI port fails.

### Corrective Action

Reboot the filer. If issue still persists then replace the xfcvi adapter.

### Syslog Message

FC-VI adapter: CPR reset initiated by port %s failed.

### Parameters

**port** (STRING): xfcvi Interconnect adapter.

## **xfcvi.qlgc.cpr.reset.unsupported**

### **Severity**

ERROR

### **Description**

This message occurs when the CPR reset is not supported by the adapter.

### **Corrective Action**

Reboot the filer. If issue still persists then replace the xfcvi adapter.

### **Syslog Message**

FC-VI adapter: CPR reset is not supported by port %s.

### **Parameters**

**port** (STRING): xfcvi Interconnect adapter.

## **xfcvi.qlgc.FrameDiscardErr**

### **Severity**

ERROR

### **Description**

This message occurs when the FC-VI adapter firmware discards a received frame because of a receive frame error.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Discarded a received frame because of a receive frame error. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**qp\_name** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.hard.reset.failed**

### **Severity**

ERROR

### **Description**

This message occurs when an FC-VI port's soft reset fails.

### Corrective Action

If CPR reset is not supported in the platform, then reboot the filer to recover from this issue. If issue persists then replace the xfcvi adapter.

### Syslog Message

FC-VI adapter: Hard reset failed at port %s.

### Parameters

**port** (STRING): xfcvi Interconnect adapter.

## xfcvi.qlgc.invalid.data

### Severity

ERROR

### Description

This message occurs when xfcvi driver receives invalid data from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid data. QP index = %d, completion status = %s(status code = %d), opcode = %s(%d), remote node's system id = %llu. %s

### Parameters

**port** (STRING): xfcvi Interconnect NIC.

**qpnum** (INT): QP index.

**status** (STRING): Completion status.

**status\_code** (INT): Completion status code.

**opcode\_str** (STRING): xfcvi opcode in string.

**opcode** (INT): xfcvi opcode.

**systemID** (LONGINT): Remote node's system Id.

**info** (STRING): Additional info.

## xfcvi.qlgc.invalid.exch.address

### Severity

ERROR

### Description

This message occurs when xfcvi driver receives a response with invalid exchange address.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid exchange address for %s. QP name = %s, QP index = %d, expected exchange address = 0x%x, received exchange address = 0x%x. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**resp\_type** (STRING): Response type.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**exp\_oxid** (INTHEX): Expected exchange address.  
**rcvd\_oxid** (INTHEX): Received exchange address.  
**info** (STRING): Additional info.

## xfcvi.qlgc.invalid.qp

### Severity

ERROR

### Description

This message occurs when xfcvi driver receives invalid QP index from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid QP. Received QP index = %d, completion status = %s(status code = %d), entry type = %s(entry code = %d). %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qpnum** (INT): QP index.  
**status** (STRING): Completion status.  
**status\_code** (INT): Completion status code.  
**entry\_type** (STRING): Request/Response entry type.  
**entry\_code** (INT): Request/Response entry code.  
**info** (STRING): Additional info.

## xfcvi.qlgc.invalid.response

### Severity

ERROR

### Description

This message occurs when xfcvi driver receives invalid response from firmware.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Port %s received invalid response. %s. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**error\_type** (STRING): Type of error.  
**info** (STRING): Additional info.



## **xfcvi.qlgc.ioErr**

### **Severity**

ERROR

### **Description**

This message occurs when the FC-VI driver receives some IO error.

### **Corrective Action**

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### **Syslog Message**

FC-VI adapter: xfcvi driver on port %s received IO error. Status = %s(status code = 0x%x), xfcvi opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu.

### **Parameters**

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code.  
**opcode\_str** (STRING): xfcvi opcode in string.  
**opcode** (INHEX): xfcvi opcode.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.

## **xfcvi.qlgc.login.timedout**

### **Severity**

NOTICE

### **Description**

This message occurs when port login request times out.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Remote port login timedout on port %s. Remote port id = 0x%x, Remote port handle = %d.  
%s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**portid** (INHEX): Remote port identifier.  
**porthandle** (INT): Remote port handle.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.ls.reject**

**Severity**

ERROR

**Description**

This message occurs when the xfcvi firmware rejects a link service.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link service = %s(code = %d), rejected on port %s. %s

**Parameters**

**ls** (STRING): Link service name.

**ls\_code** (INT): Link service code.

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

**xfcvi.qlgc.mbox.cmd.pending****Severity**

ERROR

**Description**

This message occurs when a mailbox command is pending for a long time.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Mailbox command '%s(code = 0x%x)' pending on port %s. %s

**Parameters**

**cmd** (STRING): Mailbox command.

**cmd\_code** (INTHEX): Mailbox command code.

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

**xfcvi.qlgc.mbox.cmd.timedout****Severity**

ERROR

**Description**

This message occurs when a mailbox command times out.

**Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Mailbox command '%s(code = 0x%x)' timedout on port %s. %s

## Parameters

**cmd** (STRING): Mailbox command.  
**cmd\_code** (INTHEX): Mailbox command code.  
**port** (STRING): xfcvi Interconnect NIC.  
**info** (STRING): Additional info.

## xfcvi.qlgc.no.iu.buffer

### Severity

ERROR

### Description

This message occurs when the xfcvi firmware identifies no more incoming IU buffers in the free pool.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Port %s doesn't have free IU buffer. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**info** (STRING): Additional info.

## xfcvi.qlgc.no.recv.buffer

### Severity

ERROR

### Description

This message occurs when the xfcvi firmware identifies no more receive buffers available for the incoming receive frames.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Port %s doesn't have free receive buffer. QP name = %s, QP index = %d, Message ID = 0x%x. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**msg\_id** (INTHEX): Message id of the incoming message.  
**info** (STRING): Additional info.

## xfcvi.qlgc.oodErr

### Severity

ERROR

### Description

This message occurs when the FC-VI driver receives out of order completions.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: xfcvi driver on port %s receives out of order completions. xfcvi opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu.

### Parameters

**port** (STRING): Interconnect NIC.  
**opcode\_str** (STRING): xfcvi opcode in string.  
**opcode** (INTHEX): xfcvi opcode.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system ID.

## xfcvi.qlgc.PCIErr

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter firmware encounters a PCI error because it cannot write a frame to host memory.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: PCI error while writing a frame to host memory. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qp\_name** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.port.reset**

### **Severity**

ERROR

### **Description**

This message occurs when a xfcvi port resets.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Resetting the port %s. %s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**reason** (STRING): Reason for port reset.

## **xfcvi.qlgc.qp.create.err**

### **Severity**

ERROR

### **Description**

This message occurs when QP creation fails.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Failed to create QP on port %s. Requested QP index = %d, completion status = %s(status code = %d). %s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**qpnum** (INT): QP index.  
**status** (STRING): Completion status.  
**status\_code** (INT): Completion status code.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.qp.destroy.failed**

### **Severity**

ERROR

### **Description**

This message occurs when the QP destroy request failed.

### **Corrective Action**

(None).

## Syslog Message

FC-VI adapter: QP destroy failed on port %s. QP index = %d. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.

**qpnum** (INT): QP index.

**info** (STRING): Additional info.

## xfcvi.qlgc.qpErr

### Severity

ERROR

### Description

This message occurs when the FC-VI driver receives a QP error from the Firmware. The FC-VI driver will attempt to retry the connection.

### Corrective Action

If this message repeats frequently, there may be an issue with the FC fabric or the adapter. Contact NetApp technical support.

## Syslog Message

FC-VI adapter: xfcvi driver on port %s received QP Error. QP name = %s, QP index = 0x%x, VI state = %s(state code = 0x%x), sub-error = %s(error code = 0x%x). Remote node's system id = %llu. This error has occurred %d times. %s

## Parameters

**port** (STRING): Interconnect Adapter.

**qpname** (STRING): QP name.

**qpnum** (INTHEX): QP number.

**state\_str** (STRING): QP state.

**state** (INTHEX): QP state code.

**suberror** (STRING): Sub-error.

**code** (INTHEX): Sub-error code.

**system\_id** (LONGINT): Remote node's system ID.

**errcnt** (INT): Number of times QP async error was encountered.

**info** (STRING): Additional info.

## xfcvi.qlgc.rcvd.invalid.iu.buffer

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter receives invalid information unit buffer.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

## Syslog Message

FC-VI adapter: Port %s received invalid IU buffer. Status = %s(status code = 0x%x), xfcvi opcode = %s(0x%x), QP name = %s, QP index = %d, buffer handle = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code.  
**opcode\_str** (STRING): xfcvi opcode in string.  
**opcode** (INHEX): xfcvi opcode.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**bufhdl** (INT): Buffer handle.  
**systemID** (LONGINT): Remote node's system ID.  
**info** (STRING): Additional info.

## xfcvi.qlgc.receive.error

### Severity

ERROR

### Description

This message occurs when the xfcvi firmware detects an error in the received frame.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Received bad frame on port %s. Mailbox-1 = %d, Mailbox-2 = %d. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**mbox1** (INT): Content of mailbox-1.  
**mbox2** (INT): Content of mailbox-2.  
**info** (STRING): Additional info.

## xfcvi.qlgc.received.disconnect

### Severity

NOTICE

### Description

This message occurs when xfcvi driver receives disconnect request from partner.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Disconnect request received on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## xfcvi.qlgc.received.reject

### Severity

NOTICE

### Description

This message occurs when xfcvi driver receives connection reject from partner.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Connect request rejected on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## xfcvi.qlgc.recvIUErr

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter receives some error in received information unit.

### Corrective Action

No corrective action is required. The driver internally resets the adapter if required. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: Port %s received IU error. Status = %s(status code = 0x%x), xfcvi opcode = %s(0x%x), QP name = %s, QP index = %d, buffer handle = %d, Remote node's system id = %llu.

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code.  
**opcode\_str** (STRING): xfcvi opcode in string.  
**opcode** (INHEX): xfcvi opcode.



**qpname** (STRING): QP name.  
**qpnum** (INT): QP number.  
**bufhdl** (INT): Buffer handle.  
**systemID** (LONGINT): Remote node's system ID.

## **xfcvi.qlgc.regDumped**

### **Severity**

NOTICE

### **Description**

This message occurs when the xfcvi driver saves a diagnostic dump in a file.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Diagnostic dump for port %s saved to file %s.

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**filename** (STRING): File to which the diagnostic dump was saved.

## **xfcvi.qlgc.RespFrameErr**

### **Severity**

ERROR

### **Description**

This message occurs when the adapter firmware sends a response frame with the error bit set.

### **Corrective Action**

(None).

### **Syslog Message**

FC-VI adapter: Firmware sent a response with the error bit set. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.  
**qpname** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.risc.paused**

### **Severity**

ERROR

**Description**

This message occurs when xfcvi RISC is paused.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Port %s detected a RISC pause error. %s

**Parameters**

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

**xfcvi.qlgc.rmt.link.down****Severity**

NOTICE

**Description**

This message occurs when xfcvi driver lost link to partner port.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link to partner node over port %s is down. Partner port id = 0x%x, partner node's system id = %llu. %s

**Parameters**

**port** (STRING): xfcvi Interconnect NIC.

**port\_id** (INHEX): Partner port id.

**system\_id** (LONGINT): Partner node's system id.

**info** (STRING): Additional info.

**xfcvi.qlgc.rmt.link.up****Severity**

NOTICE

**Description**

This message occurs when xfcvi driver gets link to partner port.

**Corrective Action**

(None).

**Syslog Message**

FC-VI adapter: Link to partner node over port %s is up. Partner port id = 0x%x, partner node's system id = %llu. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**port\_id** (INHEX): Partner port id.  
**system\_id** (LONGINT): Partner node's system id.  
**info** (STRING): Additional info.

## xfcvi.qlgc.sendIUError

### Severity

ERROR

### Description

This message occurs when the FC-VI adapter fails to send an information unit.

### Corrective Action

No corrective action is required. The driver internally retries. However, if you see this message frequently, there might be an issue with the FC fabric or the adapter. Contact NetApp technical support.

### Syslog Message

FC-VI adapter: Port %s sendIU status = %s(status code = 0x%x), xfcvi opcode = %s(0x%x), QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): Interconnect NIC.  
**status\_str** (STRING): Completion status in string.  
**status** (INHEX): Completion status code in hex.  
**opcode\_str** (STRING): xfcvi opcode in string.  
**opcode** (INHEX): xfcvi opcode in hex.  
**qpName** (STRING): QP name.  
**qpIndex** (INT): QP number.  
**systemID** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## xfcvi.qlgc.sent.disconnect

### Severity

NOTICE

### Description

This message occurs when xfcvi driver sends disconnect request to remote partner.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Disconnect request sent on port %s. QP name = %s, QP index = %d, Remote node's system id = %llu. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qpname** (STRING): QP name.  
**qpnum** (INT): QP index.

**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## **xfcvi.qlgc.soft.reset.failed**

### **Severity**

ERROR

### **Description**

This message occurs when an FC-VI port's soft reset fails.

### **Corrective Action**

None.

### **Syslog Message**

FC-VI adapter: Soft reset failed at port %s.

### **Parameters**

**port** (STRING): xfcvi Interconnect adapter.

## **xfcvi.qlgc.system.error**

### **Severity**

ERROR

### **Description**

This message occurs when xfcvi driver receives system error from xfcvi firmware.

### **Corrective Action**

No corrective action is required. The driver internally resets the adapter. However, if you see this message frequently, there might be an issue with the adapter/firmware. Contact NetApp technical support.

### **Syslog Message**

FC-VI adapter: System Error on port %s. %s

### **Parameters**

**port** (STRING): xfcvi Interconnect NIC.

**info** (STRING): Additional info.

## **xfcvi.qlgc.TimeOutErr**

### **Severity**

ERROR

### **Description**

This message occurs when the adapter firmware cannot send a request.

### **Corrective Action**

(None).

## Syslog Message

FC-VI adapter: Firmware could not send a request. Port %s, QP name = %s, QP index = %d, exchange ID 0x%x, address offset 0x%x, remote node's system id = %llu. %s

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**qp\_name** (STRING): QP name.  
**qp\_index** (INT): QP number.  
**exchange\_id** (INTHEX): Exchange ID.  
**address** (INTHEX): Address offset of the last good frame.  
**system\_id** (LONGINT): Remote node's system id.  
**info** (STRING): Additional info.

## xfcvi.qlgc.transfer.error

### Severity

ERROR

### Description

This message occurs when the xfcvi firmware detects a transfer error.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: %s transfer error on port %s. Affected queue number = %d. %s

## Parameters

**trans\_type** (STRING): Transfer type.  
**port** (STRING): xfcvi Interconnect NIC.  
**queue\_num** (INT): Queue number.  
**info** (STRING): Additional info.

## xfcvi.qlgc\_regdumpfailed events

### xfcvi.qlgc\_regDumpFailed

### Severity

NOTICE

### Description

This message occurs when the HA interconnect driver could not save a diagnostic dump.

### Corrective Action

(None).

## Syslog Message

FC-VI adapter: Failed to save diagnostic dump for port %s. %s.

## Parameters

**port** (STRING): xfcvi Interconnect NIC.  
**reason** (STRING): Reason for the failure.

# xfcvi.unsupported events

## xfcvi.unsupported.config

### Severity

ERROR

### Description

This message occurs when a network adapter port behavior type is set to use FC-VI on a platform that has not been validated for its use and is therefore unsupported.

### Corrective Action

Change the port behavior type to a value other than "FC-VI". Use the "system hardware unified-connect show" command to view network adaptor port information, including the current behavior type. Use the "system hardware unified-connect modify" command to change the network adapter port behavior type.

### Syslog Message

FC-VI port behavior type not supported on network adapter port %s.

## Parameters

**port** (STRING): The network adapter port whose behavior type is set to use FC-VI.

# xfcvi.update events

## xfcvi.update.link.state

### Severity

NOTICE

### Description

This message occurs when there is a change in physical link state of the Fibre Channel Virtual Interface (FC-VI) adapter.

### Corrective Action

(None).

### Syslog Message

FC-VI adapter: Physical link is %s on port %s.

## Parameters

**link\_state** (STRING): New link state.  
**port** (STRING): Interconnect NIC.

# xinetd events

## xinetd.hit events

### xinetd.hit.cps.limit

#### Severity

ERROR

#### Description

This message occurs when the rate of incoming network connections for a service exceeds the maximum allowed rate. The service is temporarily suspended and restarted after a configured wait time. Many such connection attempts can potentially disallow other users from logging in to the storage system, causing a Denial of Service (DOS) attack.

#### Corrective Action

Reduce the rate at which the network connections are issued.

#### Syslog Message

Number of incoming network connections exceeded the configured limit of %d connections per second for the service %s. This service will be stopped and restarted after %d seconds.

#### Parameters

**connections** (INT): Maximum number of network connections allowed to access the service per second.

**service** (STRING): Service invoked by xinetd.

**seconds** (INT): Number of seconds xinetd must wait before reenabling the service.

### xinetd.hit.instances.limit

#### Severity

ERROR

#### Description

This message occurs when the maximum number of SSH or Telnet requests that can be fulfilled simultaneously is reached. After the limit is reached, no new SSH or Telnet requests are accepted.

#### Corrective Action

Reduce the number of SSH or Telnet connections attempted simultaneously. The storage administrator can monitor and terminate excessive sessions, if necessary, using the security session "kill-cli" command.

#### Syslog Message

Maximum number of simultaneous incoming network connections exceeded the configured limit of %d connections for the service %s.

#### Parameters

**instances** (INT): Maximum number of SSH or Telnet requests that can be fulfilled simultaneously.

**service** (STRING): Service invoked by xinetd daemon.

## xinetd.hit.per.source.limit

### Severity

ERROR

### Description

This message occurs when the maximum number of incoming network connections per source IP address for a service is reached. New connection requests from this source IP address for the service are rejected until the existing connections are terminated. Many such connection attempts could potentially disallow other users from logging in to the storage system, causing a Denial of Service (DOS) attack.

### Corrective Action

If the remote host is retrying the network connection repeatedly, block the remote host by adding its IP address to the deny list using the "firewall policy" command.

### Syslog Message

Maximum number of network connections of %d received from the remote host %s for the service %s.

### Parameters

**limit** (INT): Maximum number of network connections for a service per source IP address.

**ipaddress** (STRING): IP address of the remote host that triggered maximum number of network connections of a service.

**service** (STRING): Service invoked by xinetd.



# xyz events

## xyz.trans events

### xyz.trans.vol.start

#### Severity

INFORMATIONAL

#### Description

This message occurs when Copy-Free Transition's volume transition job, responsible for transitioning 7-Mode volumes to clustered Data ONTAP®, starts.

#### Corrective Action

(None).

#### Syslog Message

Copy-Free Transition session:%s volume transition job for %d volumes started.

#### Parameters

**sessionId** (STRING): Unique session identification number.

**numVolumes** (INT): Number of volumes transitioning.

# zapi events

## zapi.killed events

### zapi.killed

#### Severity

NOTICE

#### Description

This message occurs when ZAPI Socket Manager (ZSM) terminates and no longer offers a connection on which a client can deliver Zephyr API (ZAPI) requests.

#### Corrective Action

(None).

#### Syslog Message

ZAPI Socket Manager (ZSM) terminated.

#### Parameters

(None).

## zapi.retired events

### zapi.retired.usageDetected

#### Severity

NOTICE

#### Description

This message occurs when ONTAPI (ZAPI) is used for ONTAP data storage management. ONTAPI (ZAPI) is approaching end of availability. Search your vendor's support site for information about the end of availability for ONTAPI and transition to the ONTAP REST API.

#### Corrective Action

Ensure that all ONTAP data storage management functions using ONTAPI (ZAPI) have transitioned to ONTAP REST API.

#### Syslog Message

ONTAPI (ZAPI), which is approaching end of availability, was last used from source address: %s, for ONTAP data storage management.

#### Parameters

**sourceAddress** (STRING): The source address from which the last ONTAPI (ZAPI) call was made.

## zapi.sf events

## **zapi.sf.up.ready**

### **Severity**

INFORMATIONAL

### **Description**

This message occurs when the M-host layer is operational and has notified the high-availability (HA) subsystem, which can then proceed with certain storage failover operations, such as giveback.

### **Corrective Action**

(None).

### **Syslog Message**

ZAPI: system node stable after startup.

### **Parameters**

(None).

## **zapi.snapshot events**

### **zapi.snapshot.success**

### **Severity**

NOTICE

### **Description**

This event occurs when an asynchronous Snapshot(tm) copy is created successfully through the Zephyr API (ZAPI) interface.

### **Corrective Action**

(None).

### **Syslog Message**

Asynchronous creation of Snapshot copy named %s in volume %s was successful. The Snapshot copy was created %d milliseconds after the request was received by ZAPI.

### **Parameters**

**snapshot** (STRING): Name of the Snapshot copy.

**volume** (STRING): Volume name in which the Snapshot copy was created.

**time** (INT): Duration, in milliseconds, that it took to create the Snapshot copy.

# zrto events

## zrto.lif events

### zrto.lif.upgrade.failed

#### Severity

ERROR

#### Description

This message occurs when ONTAP® fails to re-create a previously deleted LIF during an upgrade. The LIF must be re-created manually.

#### Corrective Action

Attempt the failed command again. If the problem persists, contact NetApp technical support.

#### Syslog Message

Failed to re-create a LIF during upgrade. Command: "network interface create %s". Reason: %s.

#### Parameters

**command\_arguments** (STRING): Arguments of the command that failed.

**error** (STRING): Error that LIF creation encountered.

# zsm events

## zsm.send events

### zsm.send.request.timeout

#### Severity

ERROR

#### Description

This message occurs when a socket timeout occurs on an intracluster Manage ONTAP® API request to an internal ONTAPI® server.

#### Corrective Action

Check your network configuration and other potential error conditions, such as the node being down or too busy. Correct the error condition and retry the operation.

#### Syslog Message

An intracluster Manage ONTAP API request timed out while waiting for an internal ONTAPI server response.

#### Parameters

(None).

## zsm.socket events

### zsm.socket.listener.create

#### Severity

NOTICE

#### Description

The event posted when the listener socket for the ZAPI Socket Manager (ZSM) service cannot be created.

#### Corrective Action

No corrective action is possible. ZSM service is not available without a reboot.

#### Syslog Message

%s initialization failure; cannot create listener socket; error code %d (%s). ZSM service is inaccessible.

#### Parameters

**zsmstring** (STRING): The string describing the type of socket that was not created; currently ZSM.

**errorcode** (INT): The error code (errno) returned from the failed ZSM listener socket create call.

**errorstring** (STRING): The character string translation of the error code (errno) returned from the failed ZSM listener socket create call.

### zsm.socket.listener.setoption

**Severity**

NOTICE

**Description**

This event is posted when the listener socket for the ZAPI Socket Manager (ZSM) service cannot be properly initialized.

**Corrective Action**

No corrective action is possible. ZSM service is not available without a reboot.

**Syslog Message**

ZSM initialization failure; cannot set listener socket option %s at level %s; error code %d (%s). ZSM service is inaccessible.

**Parameters**

**socketoption** (STRING): The socket option value that failed.

**socketlevel** (STRING): The socket option level that failed.

**errorcode** (INT): The error code (errno) returned from the failed option set call.

**errorstring** (STRING): The character string translation of the error code (errno) returned from the failed option set call.

**zsm.socket.output.event****Severity**

INFORMATIONAL

**Description**

This event occurs when the socket for the ZAPI Socket Manager (ZSM) service cannot send output data.

**Corrective Action**

This message usually means that a network error occurred or that the client closed its connection to the storage appliance. Typically, there is no corrective action that you can take. However, if the destination is available check whether you can examine the storage appliance's and client's event logs to see if networking errors were reported.

**Syslog Message**

(None).

**Parameters**

**destination** (STRING): Client that was to receive the output. If the client is disconnected and its IP address is not available, the text "Not known" is printed.

**errorcode** (INT): Error code (errno) returned from the failed write call.

**errorstring** (STRING): The character string translation of the errno returned from the failed socket write call.

# ztl events

## ztl.smap events

### ztl.smap.online.status

#### Severity

NOTICE

#### Description

This message occurs when the specified partition on a SSD-ZNS drive could not be onlined due to internal S/W or device error.

#### Corrective Action

(None).

#### Syslog Message

Disk %s partition %d (first zone=%d, number of zones=%d) online operation failed. HBA status: %d, NVMe status (type: %d, code: %d).

#### Parameters

**disk** (STRING): Disk name.

**partition** (INT): SMap Partition number.

**start\_zone** (INT): SMAP start logical zone.

**num\_zones** (INT): Number of logical zones in the SMAP instance.

**ha\_err\_status** (INT): Host Adapter error status.

**nvme\_err\_code\_type** (INT): NVME error code type.

**nvme\_err\_code** (INT): NVME error for the specific NVME error code type.

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

<https://www.netapp.com/company/legal/copyright/>

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

## Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

## Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.



## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.