



## **raid events**

### **ONTAP 9.14.1 EMS reference**

NetApp  
February 12, 2024

# Table of Contents

- raid events . . . . . 1
  - raid.aggr events . . . . . 1
  - raid.aggregate events . . . . . 2
  - raid.aggrvote events . . . . . 2
  - raid.assim events . . . . . 4
  - raid.autopart events . . . . . 21
  - raid.capacity events . . . . . 24
  - raid.carrier events . . . . . 25
  - raid.cksum events . . . . . 27
  - raid.config events . . . . . 58
  - raid.data events . . . . . 85
  - raid.debug events . . . . . 88
  - raid.disable events . . . . . 88
  - raid.disk events . . . . . 89
  - raid.diskadd events . . . . . 117
  - raid.disktoc events . . . . . 119
  - raid.enable events . . . . . 120
  - raid.encrypt events . . . . . 121
  - raid.fdr events . . . . . 121
  - raid.flexvol events . . . . . 124
  - raid.fm events . . . . . 125
  - raid.foreignaggregate events . . . . . 129
  - raid.fs events . . . . . 130
  - raid.hybrid events . . . . . 130
  - raid.initialize events . . . . . 131
  - raid.label events . . . . . 132
  - raid.labeledit events . . . . . 135
  - raid.labelmaint events . . . . . 135
  - raid.mcc events . . . . . 137
  - raid.mgrio events . . . . . 142
  - raid.mgvio events . . . . . 143
  - raid.mirror events . . . . . 143
  - raid.monitor events . . . . . 165
  - raid.multierr events . . . . . 166
  - raid.mv events . . . . . 173
  - raid.no events . . . . . 175
  - raid.ns events . . . . . 175
  - raid.nvram events . . . . . 176
  - raid.offline events . . . . . 179
  - raid.olm events . . . . . 179
  - raid.partition events . . . . . 180
  - raid.plex events . . . . . 181
  - raid.pmdopt events . . . . . 182

raid.poolsort events	184
raid.preserve events	186
raid.quiesce events	187
raid.read events	187
raid.readerr events	191
raid.recons events	197
raid.reparity events	198
raid.replay events	198
raid.rg events	200
raid.root events	261
raid.scrub events	262
raid.shared events	264
raid.sp events	266
raid.sparecore events	271
raid.spares events	271
raid.stripe events	275
raid.switchoverfail events	280
raid.switchovernodisks events	281
raid.syncmirr events	281
raid.syncmirror events	281
raid.tetris events	283
raid.tree events	289
raid.uninitialized events	289
raid.unpartition events	290
raid.unsupported events	291
raid.verify events	292
raid.vldb events	292
raid.vol events	293
raid.warn events	307

# 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 wafllron 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 wafliion 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

**bn0** (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.bn0

### 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 wafliion 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% 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% 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.

**instanceId** (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 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.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 unfail 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.multiperr.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.multierr.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** (INTEHEX): Stripe ID field of the RAID write signature on the data block.  
**gen\_cnt** (INTEHEX): Generation count field of the RAID write signature on the data block.  
**comp\_cksum** (INTEHEX): Computed RAID checksum of the disk block.  
**sto\_vbn** (LONGINT): Volume block number stored in the checksum entry.  
**wafI\_cxt0** (INTEHEX): Bytes 0-3 of WAFL® context information in the checksum.  
**wafI\_cxt1** (INTEHEX): Bytes 4-7 of WAFL context information in the checksum.  
**wafI\_cxt2** (INTEHEX): Bytes 8-11 of WAFL context information in the checksum.  
**wafI\_cxt3** (INTEHEX): Bytes 12-15 of WAFL context information in the checksum.  
**wafI\_cxt4** (INTEHEX): Bytes 16-19 of WAFL context information in the checksum.  
**wafI\_cxt5** (INTEHEX): Bytes 20-23 of WAFL context information in the checksum.  
**wafI\_cxt6** (INTEHEX): Bytes 24-27 of WAFL context information in the checksum.  
**wafI\_cxt7** (INTEHEX): 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** (INTEHEX): Stripe ID field of the RAID write signature on the diag parity block.  
**gen\_cnt1** (INTEHEX): Generation count bitmap field of the RAID write signature on the diag parity block, for data blocks 0-15.  
**gen\_cnt2** (INTEHEX): Generation count bitmap field of the RAID write signature on the diag parity block, for data blocks 16-27.  
**comp\_cksum** (INTEHEX): Computed RAID checksum of the disk block.  
**wafI\_cxt0** (INTEHEX): Bytes 0-3 of WAFL® context information in the checksum.  
**wafI\_cxt1** (INTEHEX): Bytes 4-7 of WAFL context information in the checksum.  
**wafI\_cxt2** (INTEHEX): Bytes 8-11 of WAFL context information in the checksum.  
**wafI\_cxt3** (INTEHEX): Bytes 12-15 of WAFL context information in the checksum.  
**wafI\_cxt4** (INTEHEX): Bytes 16-19 of WAFL context information in the checksum.  
**wafI\_cxt5** (INTEHEX): Bytes 20-23 of WAFL context information in the checksum.  
**wafI\_cxt6** (INTEHEX): Bytes 24-27 of WAFL context information in the checksum.  
**wafI\_cxt7** (INTEHEX): 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.

**waf\_cxt0** (INTHEX): Bytes 0-3 of WAFL® context information in the checksum.

**waf\_cxt1** (INTHEX): Bytes 4-7 of WAFL context information in the checksum.

**waf\_cxt2** (INTHEX): Bytes 8-11 of WAFL context information in the checksum.

**waf\_cxt3** (INTHEX): Bytes 12-15 of WAFL context information in the checksum.

**waf\_cxt4** (INTHEX): Bytes 16-19 of WAFL context information in the checksum.

**waf\_cxt5** (INTHEX): Bytes 20-23 of WAFL context information in the checksum.

**waf\_cxt6** (INTHEX): Bytes 24-27 of WAFL context information in the checksum.

**waf\_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 wafll 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).

## 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.